IBM Tealeaf cxImpact
Version 9 Release 0
June 12, 2014

cxImpact Administration Manual

Note
Before using this information and the product it supports, read the information in "Notices" on page 463.

This edition applies to version 9, release 0, modification 0 of IBM Tealeaf cxImpact and to all subsequent releases and modifications until otherwise indicated in new editions.

© Copyright IBM Corporation 1999, 2014.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

cxImpact Administration Manual xi

Chapter 1. cxImpact Admin Overview . . 1

What is Tealeaf cxImpact? 1
How cxImpact Works 1
Tealeaf Terms and Concepts 3
Passive Capture 3
Processing 4
Search, Replay, and Reporting 5
cxImpact Features 5
Logging in as Admin 5
Login requirements 6
Browsers 6
Usernames and Passwords. 6
Logging In 6

Chapter 2. Tealeaf System Monitoring . . 7

Identifying Issues in Tealeaf 7
User Activities 7
Portal Reporting for Tealeaf Administrators 7
External Monitoring Tools 7
Using Tealeaf Alerts to Trigger External Programs . . 8
System Status 8
Tealeaf System Status Dashboard 8
Active Status tab 9
Canister Status tab 10
DecoupleEx Status tab 10
Storage Status tab 11
Database Size Tab 12
General Usage of System Status Reports 13
Configuring a System Status Report 13
Toolbar 14
Reports 15
Canister Status 15
DecoupleEx 22
Events 26
Storage 26
Health Based Routing (HBR) 30
Database Filegroup Size 33
Database Table Size. 35
Collection Status. 37
System Statistics 39
Enabling System Statistics Hits 39
Statistics Logging 40
Generating a System Statistics Report. 40
Saving System Statistics Reports 42
User-Defined Reports 42
Pre-Configured Reports 42
Component Status 42
Hit Flow 42
Byte Flow 43
Queuing 45
CX Capture Hits. 46
CX Capture Memory 47
CX Capture Queue 47
CX CPU 49
Active Sessions 50
Canister Reporting Data 51
Indexing 52
STC Performance 53
Canister Memory 55
Canister CPU 56
Canister Statistics 57
Monitoring User Activity 57
Configuring User Activity Reports. 57
Toolbar 58
Date Range Selector 58
Options. 60
Common report fields 62
Audit Reports 62
User Activity Report 62
User Inactivity Report 67
Usage Metrics Summary Report 68
Usage Metrics by User Report 69
User Searches Report 71
Searches with Results Report 71
Sessions Returned Report. 72
Sessions Replayed Report. 73
Searches without Results Report 74
Search Effectiveness Report 74
Search Efficiency Report 75
Replay Use Report 76
Export 77
Portal Application Information Report 78
Portal Performance Tests 78
Database Information 79
Portal Component Names 79
Tealeaf Administrator 79
Time Zone Offset 80
System Properties 80
Monitoring Tealeaf Using External Tools. 81
Basic Connectivity 81
Web browser 81
Nslookup 82
Tracert 82
Capture 82
tcpdump 82
Wireshark 82
Windows Servers 82
PerfMon 82
Windows Services Control Panel 82
Process Explorer 83
Database 83
SQL Manager. 83
SQLIO 83
Replay 83
Fiddler 83
YSlow 83
HTTPwatch 83
Reporting for Tealeaf Administrators 84
Monitoring Current System Status 84

© Copyright IBM Corp. 1999, 2014

Active menu 84
System Dashboards 84
Activity Reports 85
Performance Reports 85
System Status Report 85
Checking Component Status through TMS . . . 86
Tracking System Statistics. 86
Reporting on Tealeaf Servers 86
Portal Management Logs 86
Via TMS 87
Reporting on Tealeaf Services 87
Reporting on TMS 88
TMS Logs 88
TMS in Tealeaf Status 88
Reporting on the Portal Application 88
Portal Application Information Report 88
Tealeaf Status Report 89
Reporting on Users and Groups 89
Reporting on Captured Hits 90
Hits in the PCA 90
Hits in the Transport Service or HBR 90
In the Windows Pipeline 91
After Pipeline in STC 91
Reporting on LTC 91
Reporting on Tealeaf Databases. 92
Database-Sourced Portal Reports 92
Portal Database Queries 93
Data Collector Status 94
Visitor DB Extractor Status 94
Tealeaf Database Manager 95
Reporting on Integration 95
cxConnect for Data Analysis. 95
cxVerify 96
PCA Administration Reporting 96
Network Traffic 96
PCA Application. 96
Monitoring Capture Status through TMS . . . 97
Monitoring Capture Status through System Statistics 97
Troubleshooting Tealeaf CX Issues through the Portal 97
Monitoring Active Status 97
No Current Activity 99
Troubleshooting Issues in the Pipeline 99
No Hits in the Pipeline 99
Pipeline Data Sources 100
Troubleshooting Issues in the Canister 101
Canister Health. 102
Short Term Canister 102
Long Term Canister 108
Troubleshooting Issues in Reporting 109
Identifying if the relevant services are on . . . 109
Review database activity reports 110
Check Reports for Data 110
Reviewing Data Collection Events and Logs . . 111
Troubleshooting Visitors Database Issues 112
Search 113
Activity Reports 113
Visitor Database Extractor log 113
About Page in the Portal 114
Next Steps 114

Chapter 3. Portal Management 115

Managing Tealeaf Servers 115
Managing Your Servers 116
Best Practices 116
Adding a New Server 117
Editing a Server's Properties 119
Server Toolbar 119
Deactivating a server 125
Portal Server 125
cxReveal Database Search Server 125
Replay Server 127
Proxy Server Authentication for Replay Server 130
On-Demand Privacy 130
Data Service Server 130
Visitor Report Server 131
Report Search Server 131
Canister Server 131
Canister Search Server 133
Selective Archive Server 133
Transport Server 134
TLI Server 134
TLI Files 135
cxConnect Server 136
cxVerify Server 137
Capture Application Server. 137
Configuring Browser-Based Replay 138
How BBR Works 139
How BBR accesses images 139
Replay Configuration. 140
Replay Rules 140
On-Demand Privacy 140
Authentication 140
Authentication Issues with Replay 140
Other Supported Formats 141
Configuring Search Templates 141
Search Templates 141
Template Search Fields 145
Session List Templates 155
Session List Columns 159
Building Searches via Encoded URL 163
CX User Administration 165
Users 165
New / Edit 165
Password. 168
Force Password Change 168
Clear Password History 168
Assign Groups 168
Manage NT Users 168
View Activity 168
Delete 169
Groups 169
New / Edit 169
Assign Users 171
Menu Profile 171
Administration Profile 171
Search Profile 172
Browser Replay Profile 173
Delete 174
Current Users 174
cxReveal User Administration 174
Users 175

Groups . . . 178
Current Users . . . 183
Portal Menu Bar . . . 183
cxView User Administration . . . 184
Users . . . 184
New / Edit . . . 184
Password . . . 186
Force Password Change . . . 186
Clear Password History . . . 186
Assign Groups . . . 186
Manage NT Users . . . 186
View Activity . . . 187
Delete . . . 187
Groups . . . 187
New / Edit . . . 187
Assign Users . . . 188
Menu Profile . . . 188
Administration Profile . . . 189
Delete . . . 189
Current Users . . . 189
Configuring cxView-only Users . . . 189
Expanding cxView User Permissions . . . 190
Example - Enabling Drill-Down Search for cxView Users . . . 190
CX Settings . . . 192
Do not change these settings . . . 192
Alerts/Thresholds . . . 193
Data Collector . . . 193
Data Retention . . . 201
Search . . . 204
Evaluating Search Results Limits . . . 206
Users . . . 207
Integration . . . 208
Miscellaneous . . . 211
cxReveal Settings . . . 213
Search . . . 213
Users . . . 215
cxView Settings . . . 216
Scorecards . . . 217
Users . . . 218
cxResults Settings . . . 219
Do not change these settings . . . 219
Visitor Extract . . . 219
Visitor Server . . . 222
RS Extractor Settings . . . 223
Don't Change These Settings . . . 223
Service Settings . . . 223
Field Settings . . . 225
Aliases . . . 225
Report Configuration . . . 226
Portal Announcements . . . 226
Enabling Announcements . . . 226
Configure Announcements . . . 227
New / Edit . . . 227
Up / Down . . . 227
Refresh . . . 227
Portal Logs . . . 227
Global System Logging Level . . . 228
Reporting Service, Data Collector and Query Server . . . 228
Data Service . . . 229
RS Extractor . . . 229
Session Segments . . . 230
Portal History . . . 230
Tealeaf Status Audit Log Events . . . 231
View Tealeaf Status . . . 231
Tealeaf Logs and Configuration Files . . . 231
Tealeaf Status Report . . . 231
Enabling Tealeaf Status Report . . . 231
Configuring Tealeaf Status Report . . . 231
Generating Tealeaf Status Report . . . 233
Sections of Tealeaf Status Report . . . 234
Data Aggregation and Retention . . . 241
Session Data Retention . . . 241
Data Collector Service . . . 241
Data Collection . . . 242
Data Aggregation . . . 242
cxImpact . . . 242
Performance Data . . . 243
Dimension Logs . . . 243
Managing Data Aggregation Performance . . . 244
Data Trimming . . . 245
System Data . . . 246
Reporting Data . . . 246
Tealeaf System Statistics Data . . . 246
Dimension Data Trimming . . . 246
Data Retention . . . 246
General effects of increasing data retention . . . 247
Reporting data retention . . . 247
System Processing Stats Data Retention . . . 247
User Activity Log Data Retention . . . 247
Data Retention for cxResults . . . 248
Retaining Data after Database Expiration . . . 248
Configuring Locale Settings . . . 249
Configuring Language for the Portal . . . 249
Configuring regional and language settings . . . 249
Portal Login Page . . . 249
DNS Computer Names . . . 249
Configuring Languages for User Accounts . . . 249
Configuring Time Zone . . . 251
Configuring Tealeaf Time Zone . . . 251
Time Zones for Search . . . 251
Configuring Currency . . . 252
Configuring the Canisters to Recognize Currency Identifiers . . . 252
Supported Languages . . . 253
Encoding . . . 253
Logging . . . 253
Data Service log . . . 253
About page . . . 253
Configuring Character Encoding . . . 253
Backing up the IBM Tealeaf event model . . . 254
Overview . . . 254
Limitations . . . 254
Archived Event Model Data . . . 254
Backup Data Files . . . 255
Configuring Backup Settings . . . 255
Managing Backups . . . 256
Creating On-Demand Backups . . . 256
Restoring Backups . . . 256
Deleting Backups . . . 257

Scenario - Migrating Events between Dev and Test. 257
Logging 258
Backing Up Tealeaf 258

Chapter 4. Tealeaf Management System 259
Overview of TMS 259
External Configuration Management Tools . . . 260
Accessing TMS 260
TMS Tabs. 261
Pipeline Editor 261
TMS Administration 262
TMS Reporting 262
TMS WorldView Tab 262
Selecting View 262
Applying Display Filters 263
Managing Keywords 265
Adding keywords 265
Importing and exporting templates 266
Servers View 266
Server Actions 266
TMS Master Server 266
Components View. 267
Component Actions 267
Configs View 268
Config Actions 268
Viewing and editing configurations 269
Editing configurations 270
Assigning configurations 272
Un-assigning a configuration 272
Renaming configurations 273
Revision history 273
Editing configuration filenames 275
Copying configurations 275
Encrypting configurations on disk 275
Update Servers 276
Refreshing the display 277
Tealeaf Component 277
Tealeaf global configuration settings 277
Tealeaf mail configuration settings 278
Shared configuration information. 278
Component Configuration References 280
Proactive alerts 281
Configuring proactive alerts for IBM Tealeaf components 281
Adding alert conditions 283
Editing alert conditions 284
Proactive alerts for Canister 286
Proactive alerts for data services 286
Proactive alerts for PCA 287
Patch Management with IBM Endpoint Manager 288
Applying patches on the IBM Tealeaf network 288

Chapter 5. TMS Jobs Tab 291
Overview. 291
Tasks Pane 292
Jobs Pane. 294

Chapter 6. TMS Pipeline Status Tab 297
Overview. 297
Connections 298
Pipelines 300

Chapter 7. TMS Advanced Tab 303
Overview. 303
Backup/Restore 304
Backup TMS Database 304
Restore TMS Database 304
Export TMS Config Summary 304
Discovery 305
File Management 305
Get File 305
Put File 305

Chapter 8. TMS Pipeline Editor 307
Overview. 307
Editing Pipeline Globals 308
Adding Pipelines 309
Building Pipelines 310
Editing Pipelines 310
Removing Pipelines 311
Adding Session Agents 311
Editing Session Agents 311
Removing Session Agents 312
Saving Session Agents 312
Available Session Agents 312
Identify Deprecated Session Agents 313
Pipeline Development Workflow 313
Configuring the TMS Server 314
Restarting TMS. 314
Reinitializing TMS Datastore 315
Troubleshooting 315
Tutorial - TMS 315
TMS Components 316
Install TMS 316
Servers 316
Disable a Server 317
Enable a Server. 317
Group a Set of Servers 317
Filter the WorldView to Show a Group of Servers 317
Components. 317
Check a Component's Current State 317
Start/Stop/Restart a Component 318
Add a New Component 318
Remove a Component 318
Group a Set of Components 318
Filter the WorldView to Show a Group of Components. 319
Configurations 319
Edit a Configuration 319
Revert a Configuration to a Previous Version 319
Share a Configuration with Multiple Servers 320
Make a Copy of a Configuration 320
Change the Description for a Configuration . . 320
Group a Set of Configurations. 320
Filter the WorldView to Show a Group of Configurations 320

Pipeline Editor 321
Edit Other Configuration Settings 321
Edit Pipeline Settings 321
Edit SessionAgent Settings 321
Add a Pipeline 322
Remove a Pipeline. 322
Add a SessionAgent to a Pipeline 322
Remove a SessionAgent from a Pipeline . . . 322
Move a SessionAgent within a Pipeline. . . . 322
Add a Custom SessionAgent 323
Jobs 323
Create a New Job 323
Add Tasks to a Job 323
Submit a Job 323
Cancel a Job. 324
View Job Status. 324
Make a Job Dependent on Another Job 324
Troubleshoot a Failed Job 324
Pipeline Status 324
View Pipeline Status for a Server 325
Switch to a Different Pipeline 325
Troubleshoot Pipeline Status 325

Chapter 9. Advanced Administration Topics 329

Authentication 329
Supported Authentication Methods 329
Portal (Database) Authentication 329
NT (AD - Active Directory) Authentication . . . 330
Mixed Mode 332
Web Services Authentication 333
Single Sign On (SSO) Authentication 333
Additional Required Configuration on IIS 7 for Windows Server 2008. 335
For NT Authentication 335
For non-NT Authentication. 337
Configuring Automatic User Creation in NT Authentication 338
Enabling NT Authentication 338
Assigning Tealeaf Product for New NT Users 338
Configure User Auto-Creation for Individual Products 339
Custom HTML content for the login page 339
Custom Menu Group. 340
CustomMenu Element 340
Items Element 341
Item Element 341
Divider Element 341
Character Indexing 342
Overview. 342
Alphabet.dat 342
Character Categories 343
Replaced Characters 343
Treatment of special characters before submitting a search query. 344
Scheduling Tealeaf Jobs 344
Scheduling Overnight Tealeaf Tasks 344
Suggested Scheduling of Overnight Tasks 344
Other Tasks 346
TLBackup and TLRestore 346
Coordinating CanTrim and TLBackup in All-in-One Environments 347
Running TLRestore 347
Configuring TLBackup 348
TLBackup performs the following checks on startup 349
Format of the TLBackup directories 349
Scheduling Canister Backups 350
Scheduling Daily Full Backups of the Long Term Archive 350
Restoring Tealeaf Data from TLBackup 350
Command-Line Options 351
Backing Up Tealeaf 352
Configuring Maximum Capture Sizes for Pages, Hits, and Sessions 352
Monitoring Oversized Hits 352
Monitoring for PCA 352
Monitoring for Tealeaf Pipeline 353
Configuring Max Data Size. 353
Configuring for PCA 353
Configuring for the Pipeline 355
Configuring Session Size 355
Implementing Tealeaf in a Secure Environment . . 356
Pre-Install Considerations 356
PCA Security 356
Operating System Security 356
Web Console 356
Private Keys. 357
Windows Server Security for IBM Tealeaf 358
Operating System Security 358
IIS Security 358
Securing the Platform and Portal 358
Portal Authentication Modes 358
Tealeaf Install Directory 359
User Accounts 359
Tealeaf Canister Security 361
Tealeaf Databases 361
Permissions 361
SQL Required Logins. 361
RTV Security 362
Best Practices 362
Data Security 362
Unauthorized Login Access. 362
Segment Builder Service. 363
Uploading Segment queries 363
Creating Segments for integration 363
Monitoring Segment Builder Service. 366
Configuring Segment Builder Service 366
Managing Your Tealeaf License Key 366
Required Upgrade Path 367
Transitioning to IBM 367
Release 8.7 or later 367
Release 7.2 or earlier 367
Release 8.0 to Release 8.6 367
IBM Passport Advantage 367
Acquiring Your Current Key 367
Acquiring the Tealeaf License File Generator . . . 368
Create an Account with IBM Passport Advantage 368
Using the Tealeaf License File Generator 368
Expiration 368

Portal Modules 368
Licensing Dependencies 369
Advanced Modules 370
Customer Information 371
Generating a License Key 371
For Upgrades 371
For New Installs 372
Validating a License Key 372
Applying a License Key through the Portal . . . 372
Applying the License Key through the Registry 373
When Your License Key Expires 373
Managing Static Archives 374
Overview of Static Archives 375
Local Archives 375
TLI Server Archives 376
Architecture 376
How Static Content is Stored 376
How Static Archives Are Used at Runtime. . . 377
Assess Content Types to Capture 377
Assess Impact on Data Storage and Throughput 378
Deployment Models 379
Configuration 380
1. Install Tealeaf Software 381
2. Configure a TLI Server Reference 381
3. Configure the Static Archive Pipeline . . . 382
4. Configure PCA Capture Types 382
5. Configure Replay Server Use of TLI Servers 383
Testing Static Archives 383
Actions on Static Archives 384
Static Archives Lifecycle 384
Rolling Archives 385
Merging Archives 385
Deleting Static Content from an Archive . . . 385
Loading Content from a File 386
Replay using Static Archives 387
BBR 387
RTV 387
Monitoring Static Archives 387
Viewing Static Archive Contents 388
Static Archives in System Status 388
Archive Status in Tealeaf Status 388
Logging 389
Managing TLI Files 389
Reporting 389
Reference 389

Chapter 10. Database Administration 391

Chapter 11. SQL Server authentication 393
Authentication 393
Required SQL Server login permissions. . . . 393
Installation login 393
Administrative login 394
User login 394
ssclr_login login 395
Tealeaf database security 395
Reference 395

Chapter 12. Configuring the Tealeaf Data Service 397
Connections 397
Installation 398
Database Configuration 398
Configuring Tealeaf Data Service 398
Starting and Stopping the Service 399
Logging 399
Accessing Logs 399
Log File Location 399
Changing the Log Level 399
Per-Minute Data Service Stats 400

Chapter 13. CX databases 403
System database 403
Filegroups 403
Sizing 403
Reporting database 403
Filegroups 404
Sizing 405
Statistics database 405
Filegroups 405
Sizing 405
Reports 405

Chapter 14. cxResults databases . . . 407
Session Segment database 407
Visitor Staging database 407
Visitor Reporting database 408
Reports 409

Chapter 15. Database backup strategy 411
Backup mechanisms 411
Recommended backup strategy 411
Alternate backup strategies 411
Database backup scripts 412
Backup notifications 412
Setting up database backups using the scripts 412
SQL Server notifications 413
Reference 413

Chapter 16. Database sizing 415
Setting the recovery model and logging levels for Tealeaf databases 415
Reference 415

Chapter 17. Troubleshooting Tealeaf databases 417
SQL Server configuration 417
Unable to connect to SQL Server 417
Failure to install due to model database size mismatch 417
Configuring use of SQL Server dynamic ports 418
Migrating to a new SQL Server 419
Poor SQL Server performance on 64-Bit operating systems 419
SuperSocket information: (SpnRegister): error 1355 in Windows application event log 420

SQL Server blocks access to procedure sys.sp_OAGetProperty 420
Tealeaf database configuration. 421
Unable to connect 421
During upgrade, Tealeaf Database Manager fails to upgrade the Result Set Extractor 421
Failing to create database in clustered server environment. 421
Tealeaf database performance 422
Tealeaf Database Manager hangs when upgrading the System database 422
Issues with individual databases 423
Visitor database for cxResults 423
Result Set Extractor Database for cxResults . . 423
Search database for cxReveal 424
Configuration issues 424
Data collection 424
Visitor Database Extractor 426
Unable to search for visitors 426
Troubleshooting database upgrades 426
Invalid filegroup specified error during database upgrade 426
Checklist 427

Chapter 18. Report Access Permissions 429
Scorecard Access Permissions 429
General Admin Permissions 429
Access Permissions for Individual Scorecards 429
Dashboard Access Permissions 430
General Admin Permissions 430
Access Permissions for Individual Dashboards 430
Report Access Permissions 431
General Admin Permissions 431
Access Permissions for Individual Reports. . . 431
Top Mover Access Permissions 432
Session Segment Access Permissions. 432
General Admin Permissions 432
Access Permissions for Individual Segments . . 433

Chapter 19. Managing User Agents 435
Overview of User Agents in Tealeaf 435
How user agent information is tracked 435
User agent strings by client type 436
How user agent information is stored 437
Enabling extended user agent detection 439
Mobile User Agents and bots 439
Public Standards 439
User Agent Tools 440
User Agent Events. 440
Tutorial 440

Chapter 20. Configuring User Agent Detection 441
Types of Files 441
`WURFL.csv` 441
`browscap.csv` 442
`UserCap.csv` 442
`UserSupplement.csv` 443
Reviewing Configuration Files. 443
Reviewing configuration information 445
Editing or creating user agent entries 445
Default User Agent Fields 447
Enable extended user agent detection 448

Chapter 21. Configuring User Agent Events 449
User Agent Evaluation Logic 449
User Agent Event Objects 449
Locating a Missing User Agent 449
Adding Values for Reporting 450

Chapter 22. User Agent Tools 451
Converting WURFL 451
Creating User Agent Event Value Lists 452
Reviewing configuration files 452

Chapter 23. bot Detection 453
Enabling bot Detection in the CX Platform . . . 453
Detecting bots 453
Masked bots. 454
bot-Related Data Objects 454
Bot (T/F) System Session Attribute 454
bot Detection in RTV 454
Traffic Type Dimension 454
Searching for bot sessions 455
Searching in RTV 455
Searching through the Portal 455
Reporting on bots 455

Chapter 24. Tealeaf Reference Information 457
Tealeaf File Extensions 457
Tealeaf Proprietary Formats 457
Other Common Formats in Use 458
Tealeaf Encryption Standards 458
Key Characteristics 458
Encryption by Component 458
Reference. 459

Chapter 25. IBM Tealeaf documentation and help. 461

Notices 463
Trademarks 465
Privacy Policy Considerations 465

cxImpact Administration Manual

The IBM Tealeaf cxImpact Administration Manual describes how to administer the IBM Tealeaf CX system, including user administration, Portal management, component settings, and more. Use the links below to access specific topics in the manual.

© Copyright IBM Corp. 1999, 2014

Chapter 1. cxImpact Admin Overview

This section is an overview to the Tealeaf® system and its components.

What is Tealeaf cxImpact?

IBM Tealeaf cxImpact is customer experience management (CEM) software that performs the following functions:

- Captures each interaction between your visitors and your website.
- Processes the low-level interaction detail to provide alerts, reports, and analysis, plus playback of those interactions.
- Shows you what happened during each visitor's experience with your site.

IBM Tealeaf cxImpact enables website owners to:

- Ensure a good experience for your visitors
- Protect revenue streams
- Reduce application support costs

Businesses are relying on Web applications to deliver business-critical services. IBM Tealeaf cxImpact addresses the challenge of managing high-volume dynamic websites by providing visibility into the production state of Web applications and allowing IT operations to ascertain whether they are functioning correctly.

IBM Tealeaf cxImpact provides a single data source that captures an audit trail of what each visitor did and saw. Tealeaf users can then analyze, isolate, replicate, and communicate problems.

How cxImpact Works

Tealeaf software running on computers in your data center is configured to see every byte of data exchanged between your web servers and the browsers on your visitors' computers. Tealeaf passively captures the bi-directional data stream and forwards a copy of each relevant packet of that data to a server running the Tealeaf system.

The figure below shows a typical IBM Tealeaf cxImpact setup with three types of Tealeaf servers.

© Copyright IBM Corp. 1999, 2014

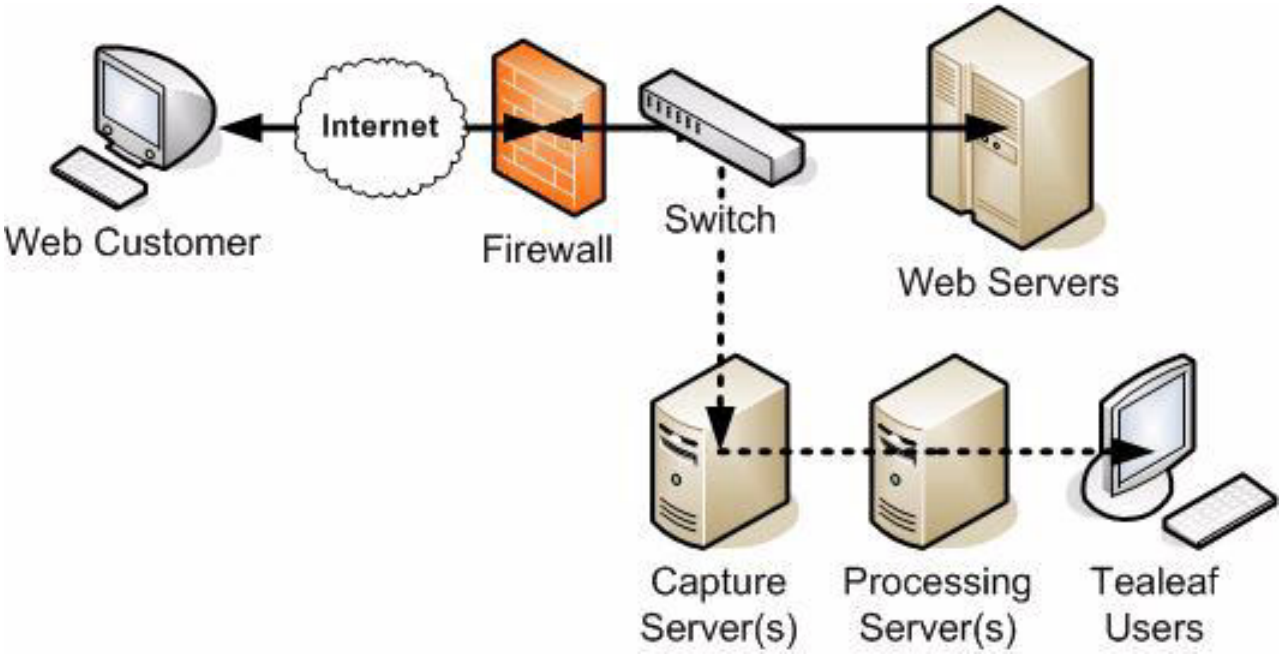

Figure 1. Tealeaf Server Network

- **Capture Server** - Tealeaf's Passive Capture software reassembles TCP/IP packets into HTTP requests and responses for each exchange between visitor and web application. Optionally, uninteresting data can be dropped, and sensitive data can be deleted. The captured combinations of requests and responses (called **hits**) are then forwarded to another server running the Tealeaf processing software.
- **Processing Server** (also called **Canister Server**) - Individual hits sent from the Capture Server are grouped together into a Tealeaf **session**, which is assembled hit-by-hit to include all web page interactions between a specific visitor and your website. The data contained in this session is scanned for keywords and codes that you have defined. When matches are found, a record of this **event** is stored for additional processing, including generation of reports and alerts. When the session is complete, the session is indexed for search and written to disk.
- **Tealeaf Portal** (also called **Report Server**) - This server's Web-based interface allows data analysts, business owners, IT staff, and administrative users to view the status of interactions with your website, with aggregate data reports that display what visitors are doing on your site. Tealeaf users can search the saved visitor sessions and analyze them for common attributes, such as order completion. Individual visitors' sessions can be replayed with the actual data that was sent and received, so a high-fidelity recreation of the visitor's experience can be reviewed.

Tealeaf Terms and Concepts

The figure below illustrates Tealeaf's data structure at a lower level.

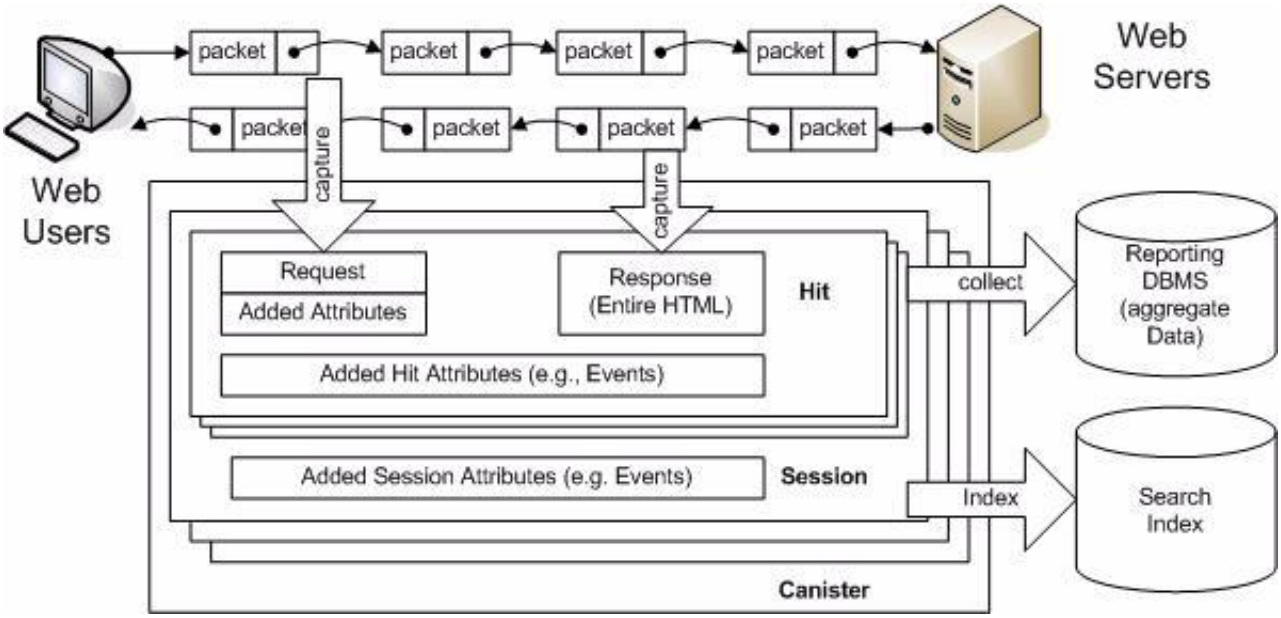

Figure 2. Tealeaf Data Structure

Passive Capture

Tealeaf's Passive Capture software runs on one or more dedicated **capture servers**.

The following terms apply to the capture process:

- **Switch** - The switch is a hardware device that routes all incoming and outgoing data packets between your visitors' computers and your web servers. Typically, switches are configured using a hardware option called a http://www.cisco.com/warp/public/473/41.html, which delivers a copy of every HTTP packet to the capture server.
- **Packet** - The TCP/IP protocol organizes interaction between computers into packets. An individual Web page may be broken down into many packets, each transmitted individually between computers. The capture server typically monitors millions of packets traveling nearly simultaneously between your Web servers and visitors' computers. These packets may arrive in any order and sometimes must be retransmitted. The capture server can be configured to ignore packets that are not of interest, such as email messages or packets sent to IP addresses of servers not hosting the website.
- **Request** - The HTTP protocol defines a request as a message requesting a response from one computer to another. The capture server collects all HTTP data to recreate the request and response traffic.
- **Response** - A response is the return message to a computer, which has made a request. After capturing a request, the capture server then processes and assembles packets in search of the response to it.
- **Hit** - A hit is defined as a request and the corresponding response to it. After the hit has been collected, the Passive Capture software can scan the data to see if the hit is of interest. For example, images that appear on every web page are not very interesting and may be discarded. Also, sensitive information such as

usernames, passwords, and credit card numbers can be deleted. After removing unwanted data, the Capture software securely transmits the hit data to the Processing Server.

- **SSL** - Many website interactions are encrypted to protect the data from being read or manipulated by third parties. The Capture software has to decrypt the data in order to match requests and responses. Typically, the Capture software is configured to re-encrypt the software using SSL for transmission to the processing servers.

Processing

The Tealeaf processing software runs on one or more dedicated **processing servers**, sometimes called **canister servers**. These servers are typically behind the company firewall and are not visible to the Internet. They accept connections from capture servers and receive hits for processing.

Each hit includes information that maps it to a specific visitor. This information is used to group the hits into a **session**. The following terms apply to Tealeaf processing:

- **Session** - As each hit is received, it is grouped with the other hits that apply to the same visitor's current interaction with your website. The processing server collects these hits until no more hits are received for the visitor or until a configured limit is reached. For example, if no more hits are received over an interval of 15 minutes, then the session may be considered complete. If a visitor performs an extended session, it may be necessary to end that session and to start a new one due time or memory constraints. While hits are arriving and being added to a session, the session is considered active. Tealeaf allows you to search for an individual session while it is still active and to view pages and interactions that have just occurred.
- **Short Term Canister** - Sessions containing hits are stored in an area of volatile memory on the processing server. As each hit arrives, it is added to a new or existing session.
- **Event** - As each hit is added to a session in the short-term canister, a series of events can be applied. An event is defined as a combination of a **condition** and an **action**.
  - The condition can be noteworthy information in the request or response code. For example, a match can be a 404 status code in the response or a particular text string such as "Purchase Confirmed."
  - When a condition evaluates to true, the action associated with the condition is performed. This action may be to increment a counter or to record a value. Event-related data is written to the short-term canister and may be aggregated into the report server. Some events are specific to a hit, while others can only be processed with the session. When a session ends and is written to the long-term canister, events that are associated with an entire session are evaluated. An event can be triggered off multiple conditions.
- **Long Term Canister** - As each session ends, either due to visitor action, lack of further hits within a timeout period, or insufficient memory to hold a longer session, it is encrypted for security and written to non-volatile storage on a hard drive. This area is called the **long term canister**.
- **Index** - As each session is written to the long term canister, the data within it is indexed for later searching. In most deployments, a selection of the most important data is indexed. After indexing, Tealeaf users can search on specific values, e.g. all sessions where the username field contained the value "smith."

- **Alert** - An alert is an action, such as sending an email, that should be taken when a predefined condition occurs. This may be the execution of an event or the crossing of a threshold, such as exceeding a total number of hits within a period of time.

Search, Replay, and Reporting

After the session data has been processed, it is available to Tealeaf users to monitor the health of their business and solve problems. Tealeaf includes several tools for these purposes.

- **IBM Tealeaf CX Portal** - This Web-based portal is the primary user interface to the Tealeaf system. It provides a wide variety of options for creating and viewing reports, including **dashboards**, **scorecards**, and other ad-hoc reports. It also provides a user interface to search for both active and completed sessions, a means to review and analyze data across search results, a means to replay a user session within a browser, and for administrative users, controls to manage the Tealeaf system. See "cxImpact User Manual" in the *IBM Tealeaf cxImpact User Manual*.
- **Search** - The indexes created by the processing server can be searched through the Portal. It is also possible to search for text strings in currently active sessions in the short term canister. But the search for completed sessions in the archive is much more powerful; Tealeaf users can search for specific types of information, such as form field values, that are not yet indexed in active sessions. See "Searching Session Data" in the *IBM Tealeaf cxImpact User Manual*.
- **Replay** - The ability to replay a session as a Web visitor experienced it is unique to Tealeaf. Replay is a powerful tool for understanding the customer experience. When you search for sessions through the Portal, you see a list of sessions matching your search criteria. You can select one of these sessions and launch a replay of the session. During replay, you see the actual Web pages server to the visitor, the links the visitor clicked on, and the form fields the visitor edited. Also available is the HTTP request data that was sent from the visitor's browser and the actual HTML returned by the Web server. See "CX Browser Based Replay" in the *IBM Tealeaf cxImpact User Manual*.
- **IBM Tealeaf CX RealiTea Viewer** - In addition to the browser user interface for search and replay, Tealeaf also supplies a Windows desktop application called the **IBM Tealeaf CX RealiTea Viewer** (RTV) that you can install on your personal computer. This application includes additional options for configuring playback on the system. See "RealiTea Viewer (RTV) User Manual" in the *IBM Tealeaf RealiTea Viewer User Manual*.

cxImpact Features

For more information on the Portal features enabled IBM Tealeaf cxImpact, see "Portal Navigation" in the *IBM Tealeaf cxImpact User Manual*.

Logging in as Admin

Note:

1. If you are unable to login or are experiencing IIS-related permission issues, you may be able to rectify the problem by re-installing the IBM TealeafCX Portal application or web service. See "Tealeaf Web Application Installation Utility" in the *IBM Tealeaf CX Installation Manual*.
2. If you have not done so already, it's recommended that you change the default administrator account passwords as soon as possible.

- See “CX User Administration” on page 165.
- See “cxView User Administration” on page 184.
- See "cxReveal User Administration" in the *IBM Tealeaf cxReveal Administration Manual*.

Login requirements

This section describes the login requirements.

Browsers

The following browsers are supported for use with the Tealeaf Portal.

Note: The Tealeaf Portal does not support access from mobile browsers.

Note: A different set of browsers and versions is supported for connecting to the Web Console of the IBM Tealeaf CX Passive Capture Application. See "Passive Capture Configuration via Web Console" in the *IBM Tealeaf Passive Capture Application Manual*.

- Internet Explorer 7 or later
- Mozilla Firefox 9 or later
- Chrome 17 or later

Usernames and Passwords

- Tealeaf usernames are not case-sensitive. `admin` and `ADMIN` both work.
- Passwords are case-sensitive. `myPassword` and `MYPASSWORD` are different.
  - Passwords may be up to 32 characters in length.

Logging In

1. Launch a supported browser.
2. Enter the following URL into your browser:

   `http://<servername>/portal`

   where: `<servername>` is the hostname of the machine on which the IBM Tealeaf CX Portal has been installed by your Tealeaf administrator.
3. Enter your login name and password.
   - If Tealeaf is configured to use the Windows Active Directory system (NT authentication), your login is the same as the name and password that you use for Windows. For more information on authentication mechanisms for the Tealeaf Portal, see “Authentication” on page 329.
4. Click **Login**.
5. The default start page for your primary user group is displayed.
   - For more information on configuring the default start page for each group, see “CX User Administration” on page 165.

Chapter 2. Tealeaf System Monitoring

Tealeaf provides several mechanisms for monitoring the health of the Tealeaf system and Tealeaf user activities. This section provides access to the various tools available in the Tealeaf system for monitoring it and its users.

Identifying Issues in Tealeaf

Through the Portal, you can troubleshoot a variety of issues that may be impacting the capture, evaluation, indexing, storage, and replay of session data.

- For a step-by-step checklist of how to use the Portal to assess and troubleshoot the IBM Tealeaf CX system, see "Troubleshooting Tealeaf CX Issues through the Portal" on page 97.

User Activities

The Tealeaf Portal captures and reports on a variety of user activities, including logins, logouts, and usage of Portal features.

- See "Monitoring User Activity" on page 57.

Portal Reporting for Tealeaf Administrators

Through the Portal, Tealeaf administrators can access a wealth of information on the current and historical performance of the IBM Tealeaf CX system as a whole, as well as individual components. Information includes a variety of useful reports, statistical information generated by the system itself, and logs for each Tealeaf component.

- See "Reporting for Tealeaf Administrators" on page 84.

The following Portal reporting features provide critical information to Tealeaf administrators who are maintaining the system or diagnosing issues.

1. **"System Status" on page 8:** These reports provide data on current activities in the Windows pipeline, Short Term Canister, Long Term Canister, and storage devices. See "System Status" on page 8.
2. **"System Statistics" on page 39:** IBM Tealeaf CX can generate statistical hits on the DecoupleEx session agent, IBM Tealeaf CX Passive Capture Application, and Canisters, which can be reviewed through the Portal. See "System Statistics" on page 39.
3. **"Tealeaf Status Report" on page 231:** The Tealeaf Status Report can be generated to provide detailed information on the status of all Tealeaf components. See "Tealeaf Status Report" on page 231.
4. **"Portal Application Information Report" on page 78:** Available through the Portal menu, this report provides status information and a set of functional tests for the Portal application. Depending on whether you are a Tealeaf user or administrator, the report delivers different content. See "Portal Application Information Report" on page 78.

External Monitoring Tools

External tools can be applied to monitor the Tealeaf solution. See "Monitoring Tealeaf Using External Tools" on page 81.

© Copyright IBM Corp. 1999, 2014

Using Tealeaf Alerts to Trigger External Programs

You can configure event-based alerts to launch external command-line programs, which can be used to provide additional monitoring and reporting capabilities.

- See "TEM Alerts Tab" in the *IBM Tealeaf Event Manager Manual*

System Status

Through the Portal, you can review System Status reports, which provide updated information on Tealeaf server and component activities.

- Summary versions of these reports are available in the Tealeaf System Status dashboard. To see the dashboard, select either of the following:
  1. **Tealeaf** > **System Status** > **Status Summary**
  2. **Dashboards** > **Tealeaf System Status** See “Tealeaf System Status Dashboard.”

Note: Depending on your account permissions, you may not have access to either or both of the above menu items. Account permissions can be configured in the Menu Profile. See “CX User Administration” on page 165.

To review the detailed System Status reports, select **Tealeaf** > **System Status**. Then, choose the report to view.

Note: You must have the appropriate permissions to access the Tealeaf menu.

For more information about System Status, see "Tealeaf System Monitoring" in *IBM Tealeaf cxImpact Administration Manual*.

Tealeaf System Status Dashboard

The Tealeaf System Status Dashboard provides visibility into the basic status of the canister, DecoupleEx session agent, and storage usage. Additionally, you can monitor the current status of events and page and session distributions.

To access the dashboard, select **Dashboards** > **Tealeaf System Status**.

Note: This dashboard is available in read-only format to all Tealeaf licensees of the IBM Tealeaf cxImpact product. Configuring the dashboard requires the IBM Tealeaf cxView license. See "Configuring Dashboards" in the *IBM Tealeaf cxView User Manual*.

Each dashboard tab is composed of multiple components.

- Click a dashboard component to see the date and time when it was last updated.

Each component contains the following commands in its toolbar:

Table 1. Tealeaf System Status Dashboard

Command	Name	Description
	Refresh	Refresh the component's data.

Table 1. Tealeaf System Status Dashboard (continued)

Command	Name	Description
	Edit	Edit the dashboard component. See "Configuring Dashboards" in the *IBM Tealeaf cxView User Manual*.
	Details	View the detail report for the component. Report details are described below in this section.

Active Status tab

The Active Status tab summarizes distributions for active sessions, single-page sessions, hits, and events.

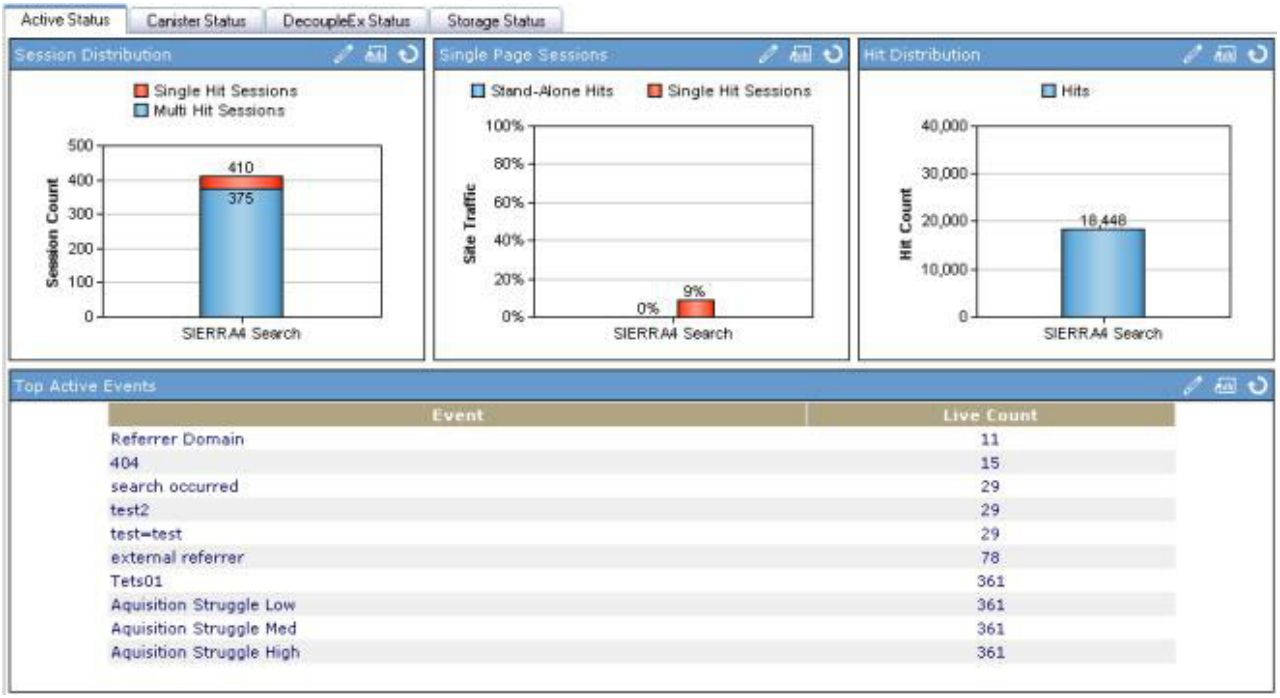

Figure 3. Tealeaf System Status dashboard - Active Status tab

Component
Description

Session Distribution
Distribution of single-hit and multi-hit sessions among the sessions stored in the Short Term Canister for all Tealeaf servers.

Single Page Sessions
Single-hit sessions and stand-alone hits as a percentage of current traffic across all servers.

Hit Distribution
The hit count among current sessions.

Top Active Events
The live counts for the most active events among current sessions

Canister Status tab

The Canister status tab identifies the current status of the Short Term Canisters across all Tealeaf Processing Servers. Each dashboard component provides insight into an important metric in identifying overload conditions.

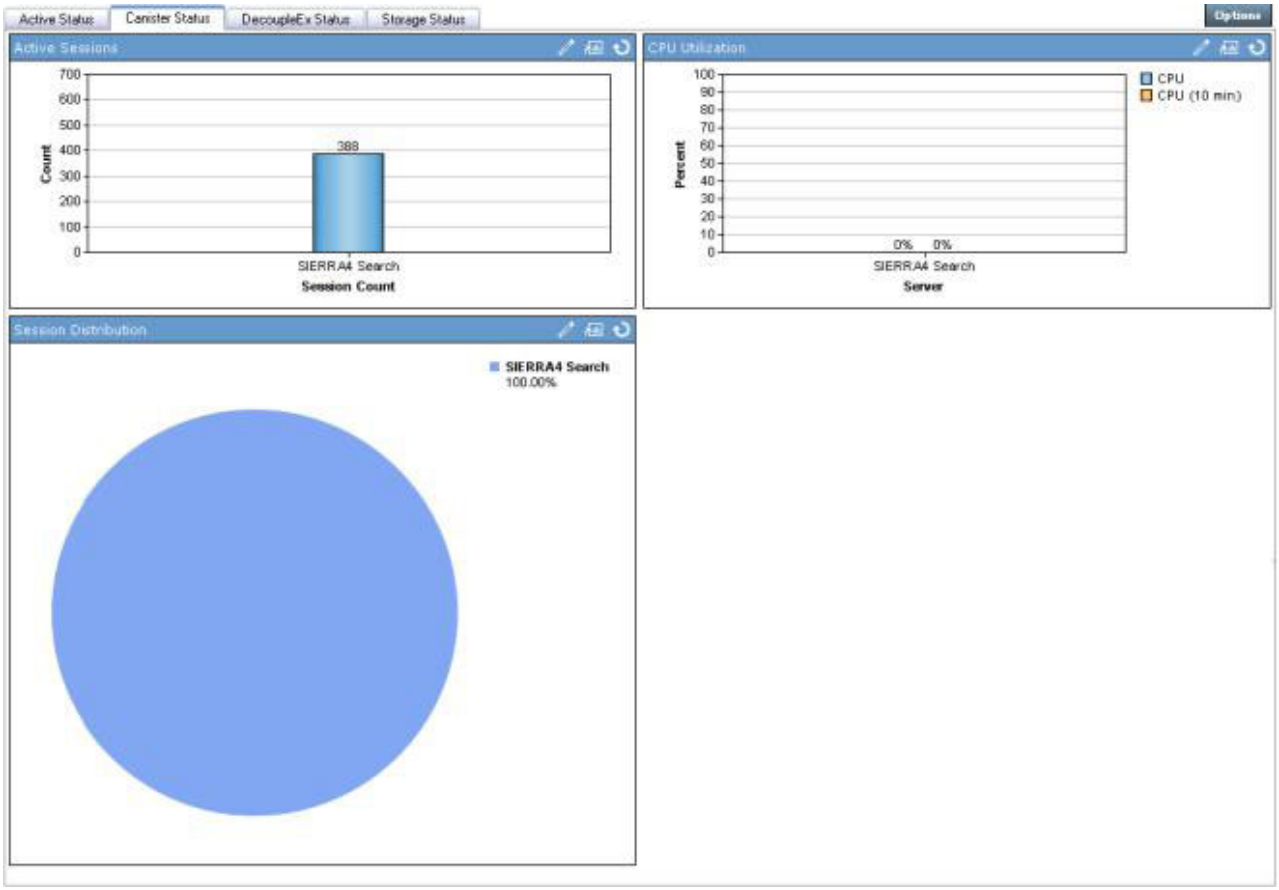

Figure 4. Tealeaf System Status dashboard - Canister Status tab

Component
Description

Active Sessions
Number of Active sessions in the canister

CPU Utilization
CPU usage now and ten minutes ago across all Tealeaf Processing Servers

Session Distribution
The distribution of current sessions across all Tealeaf Processing Servers

DecoupleEx Status tab

This dashboard displays key metrics on how the DecoupleEx session agent is performing in the processing of hits through the Tealeaf pipeline

- This report is populated from servers whose pipelines include the DecoupleEx session agent. See "Extended Decoupler Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

Figure 5. Tealeaf System Status dashboard - DecoupleEx Status tab

Component
Description

Hits/Sec Rate Overview
The rate of hits being submitted to the Short Term Canister (In), processed by it (Canister), and written to memory (Out) per second

Queue Overview
The number of hits queued in memory and on disk for processing by the STC

Storage Status tab

The Storage Status tab provides overviews of current storage usage across all Canisters in the system.

- **LSSN** files are the internal files used to store session data on the Long Term Canister disk.

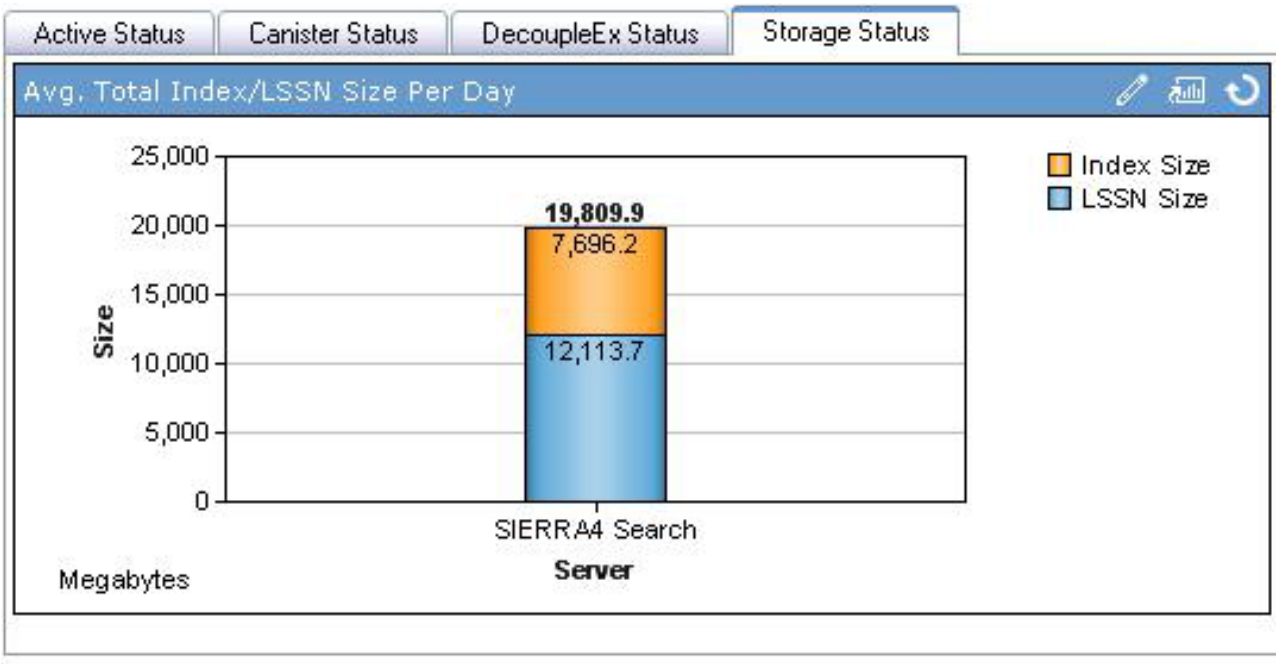

Figure 6. Tealeaf System Status dashboard - Storage Status tab

Component
Description

Avg. Total Index/LSSN Size Per Day
The average total size of the session indexes and the LSSN files per day in megabytes

Database Size Tab

The Database Size tab indicates the size and growth trends for the Tealeaf databases and table groups.

Note: If the report is indicating that one or more databases are changing size at an abnormal rate, you should investigate the source of the issue. See "Troubleshooting - Tealeaf Databases" in the *IBM Tealeaf Troubleshooting Guide*.

For more information on Tealeaf databases, see "Tealeaf Databases Guide" in the *IBM Tealeaf Databases Guide*.

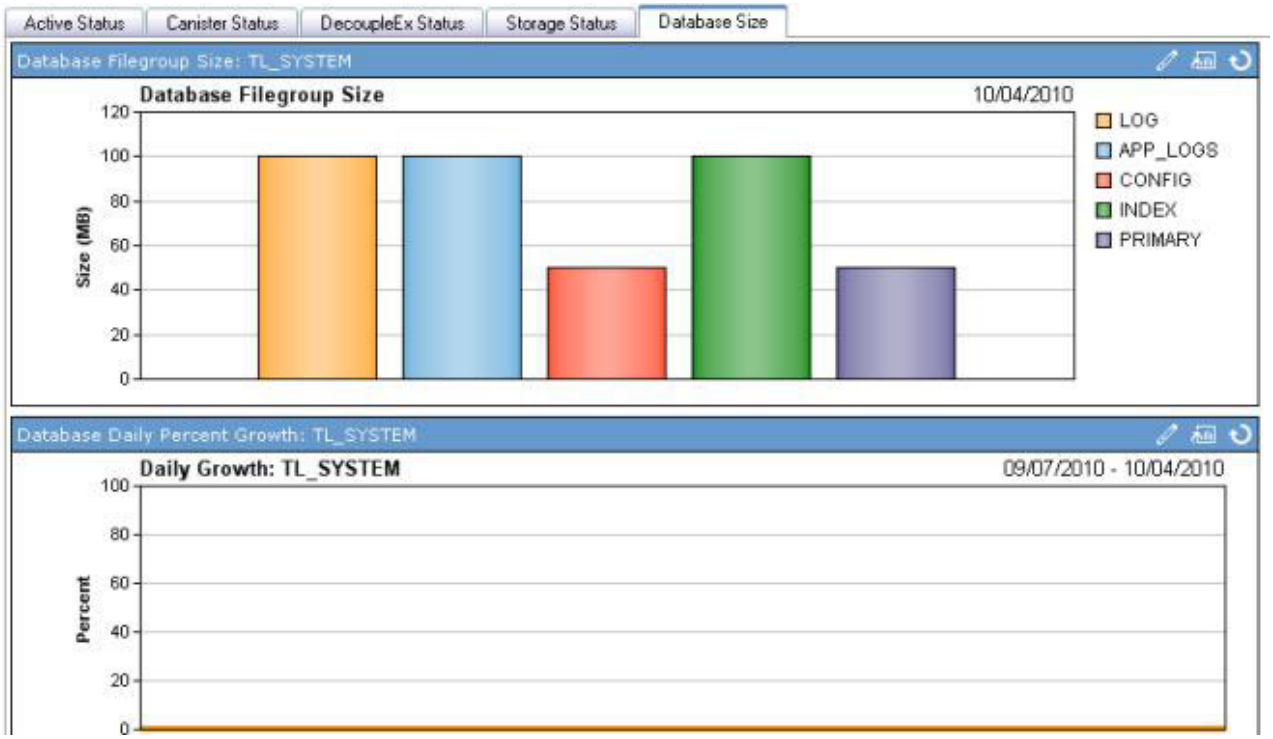

Figure 7. Tealeaf System Status dashboard - Database Size tab

Component
Description

Database Filegroup Size: <database>
Displays the current size in megabytes of the filegroups in the currently selected database.

Database Daily Percentage Growth: <database>
Displays the percentage growth day-over-day for the currently selected database over the configured reporting period.

Database Table Row Count
For each data group, you can review the number of rows currently in the tables of the group, based on the last Data Collector run.

Database Table Percent Growth: <data_group>
Displays the growth day-over-day for each selected data group over the configured reporting period.

General Usage of System Status Reports

Detailed system status reports are available through the Tealeaf Portal.

- To access one of the following reports, select **Tealeaf** > **System Status** and one of the following selections.

Status reports are available for the following components of the system:

- **"Canister Status" on page 15** - Use the Canister report to review activities of all Canisters or individual Canisters in the Tealeaf environment. This report provides current information on the Short Term Canister and Long Term Canister and can be used to debug memory allocation and processing throughput issues.
- **"DecoupleEx" on page 22** - This report provides status information on the Tealeaf Transport Service and the activities of the Decoupler, which manages traffic flow and throttles whether hits are delivered to memory or disk in the STC.
- **"Storage" on page 26** - The Storage report provides information on the storage requirements for session data and related indexes on the available canisters.
- **"Health Based Routing (HBR)" on page 30** - If Health Based Routing is deployed in your environment, this report contains information on how HBR is routing session data and the current status of the Canisters that it serves.
- **"Database Filegroup Size" on page 33** - This report is useful for monitoring the current size and growth of database filegroups from each database over a configurable date range.
- **"Database Table Size" on page 35** - This report is useful for monitoring the current size and growth of database tables over a configurable date range.
- **"Collection Status" on page 37** - Review date and time information on the most recent data collection runs for Tealeaf servers.

Configuring a System Status Report

For any selected report, you can display data in summary form or from selected servers.

1. In the Portal menu, select **Tealeaf** > **System Status**.
2. Then select one of the following reports.
   - "Canister Status" on page 15
   - "DecoupleEx" on page 22
   - "Storage" on page 26
   - "Health Based Routing (HBR)" on page 30
   - "Database Filegroup Size" on page 33
   - "Database Table Size" on page 35
   - "Collection Status" on page 37

3. To display the report for a specific server, select the server from the drop-down.

   You may also change the report by clicking the Open ( ) icon in the toolbar. See “Toolbar.”

Toolbar

When you select one of the individual reports to display, a toolbar of available commands is displayed at the top of the screen:

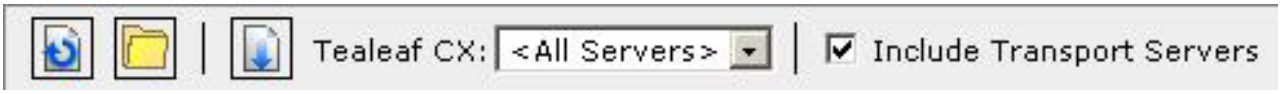

Figure 8. System Status toolbar

Command

Description

(Database Filegroup and Database Table only) For database reports, you can use the Calendar tool to select a date or range of dates to display in the report.

- To specify a range of dates in the Calendar, select the first date. Then, press SHIFT and select the second date.

Refresh the data in the report.

Open a different report. Select the report from the drop-down menu.

Export and optionally email the report data to Excel or PDF format. See “Export” on page 15.

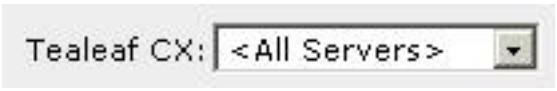

To display the report for an individual server, select the server from the drop-down.

(DecoupleEx only) To include the status of Transport Servers, which include HBR machines, select the Include Transport Servers checkbox.

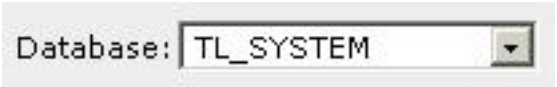

(Database Filegroup only) Select the database from which to display filegroup size information. See “Database Filegroup Size” on page 33.

(Database Table only) Select the group of database tables from which to display size information. See “Database Table Size” on page 35.

Export

To export a displayed report, click the Export ( ) icon in the toolbar. Select the type of export:

Note: Data export is not available for Canister, DecoupleEx, Storage and HBR reports.

- **Excel** - Export the report data in XML format for use in Microsoft Excel.
- **PDF** - Generate the report in PDF format.
- **Email** - Generate the report in either format and send it via email:
  1. Select the report format: PDF or Excel.
  2. Enter an optional message.
  3. Enter recipient email addresses or aliases, separated by commas.
  4. To send the report, click **Send**. The report is exported in the selected format and emailed to the specified list of addresses.

Reports

System status reports provide an easy method for monitoring the status of the Tealeaf system and its individual hardware and software components.

- For more information on active events, see "Monitoring Active Events" in the *IBM Tealeaf cxImpact User Manual*.
- For historical event counts, see "Event Activity" in the *IBM Tealeaf Reporting Guide*.

Canister Status

This report provides time-sensitive details from the Short Term Canister including number of active sessions, table sizes, CPU and memory utilization, and canister trim statistics.

- For more information on tuning the Short Term Canister, see "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

All servers

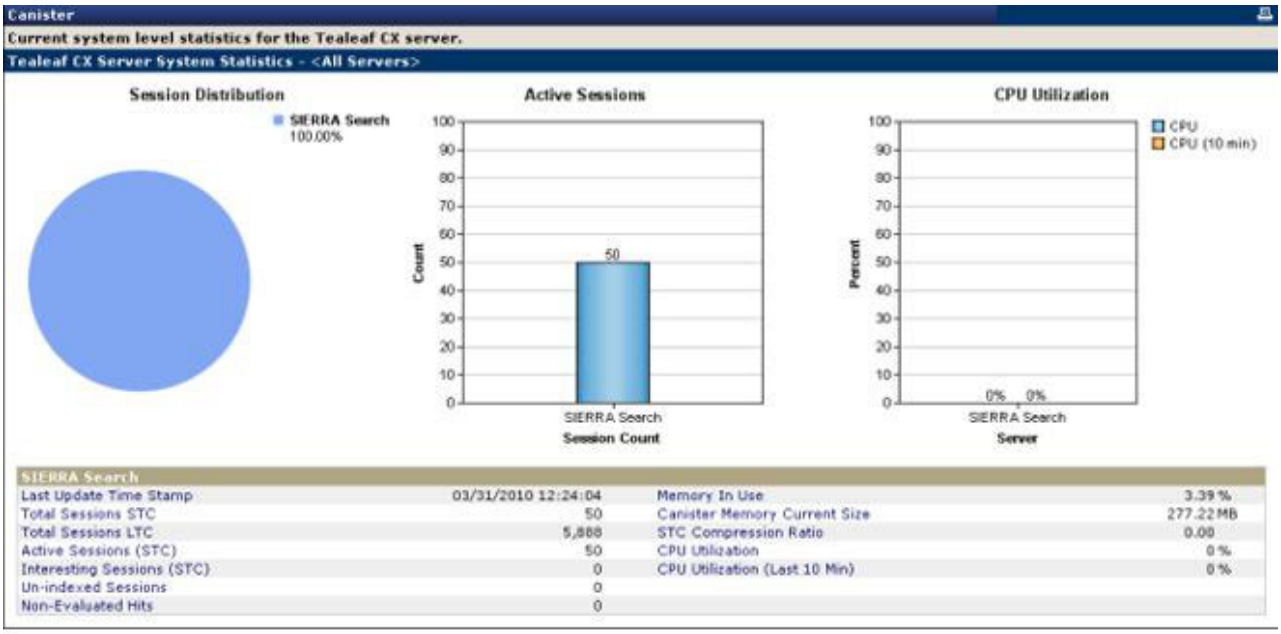

Figure 9. Canister Status report - All Servers

Field Description

Last Update Time Stamp
A timestamp indicating the last time this report was updated.

Total Sessions STC
The total number of sessions in the Short Term Canister.

Total Sessions LTC
The total number of sessions in the Long Term Canister.

Active Sessions (STC)
The total number of active sessions in the Short Term Canister.

Un-indexed Sessions
The total number of sessions waiting for indexing in the Long Term Canister.

Non-Evaluated Hits
The number of hits in the Short Term Canister that have not been evaluated.

Memory In Use
The amount of memory in use by the Short Term Canister.

Canister Memory Current Size
The current size in RAM of the Short Term Canister.

STC Compression Ratio
The rate of compression for the Short Term Canister.

CPU Utilization
The percentage rate of CPU utilization in the Canister over the reporting period.

CPU Utilization (Last 10 min)
The percentage rate of CPU utilization in the Canister over the preceding ten minutes.

Single server

When a single Canister server is selected, the following report is displayed.

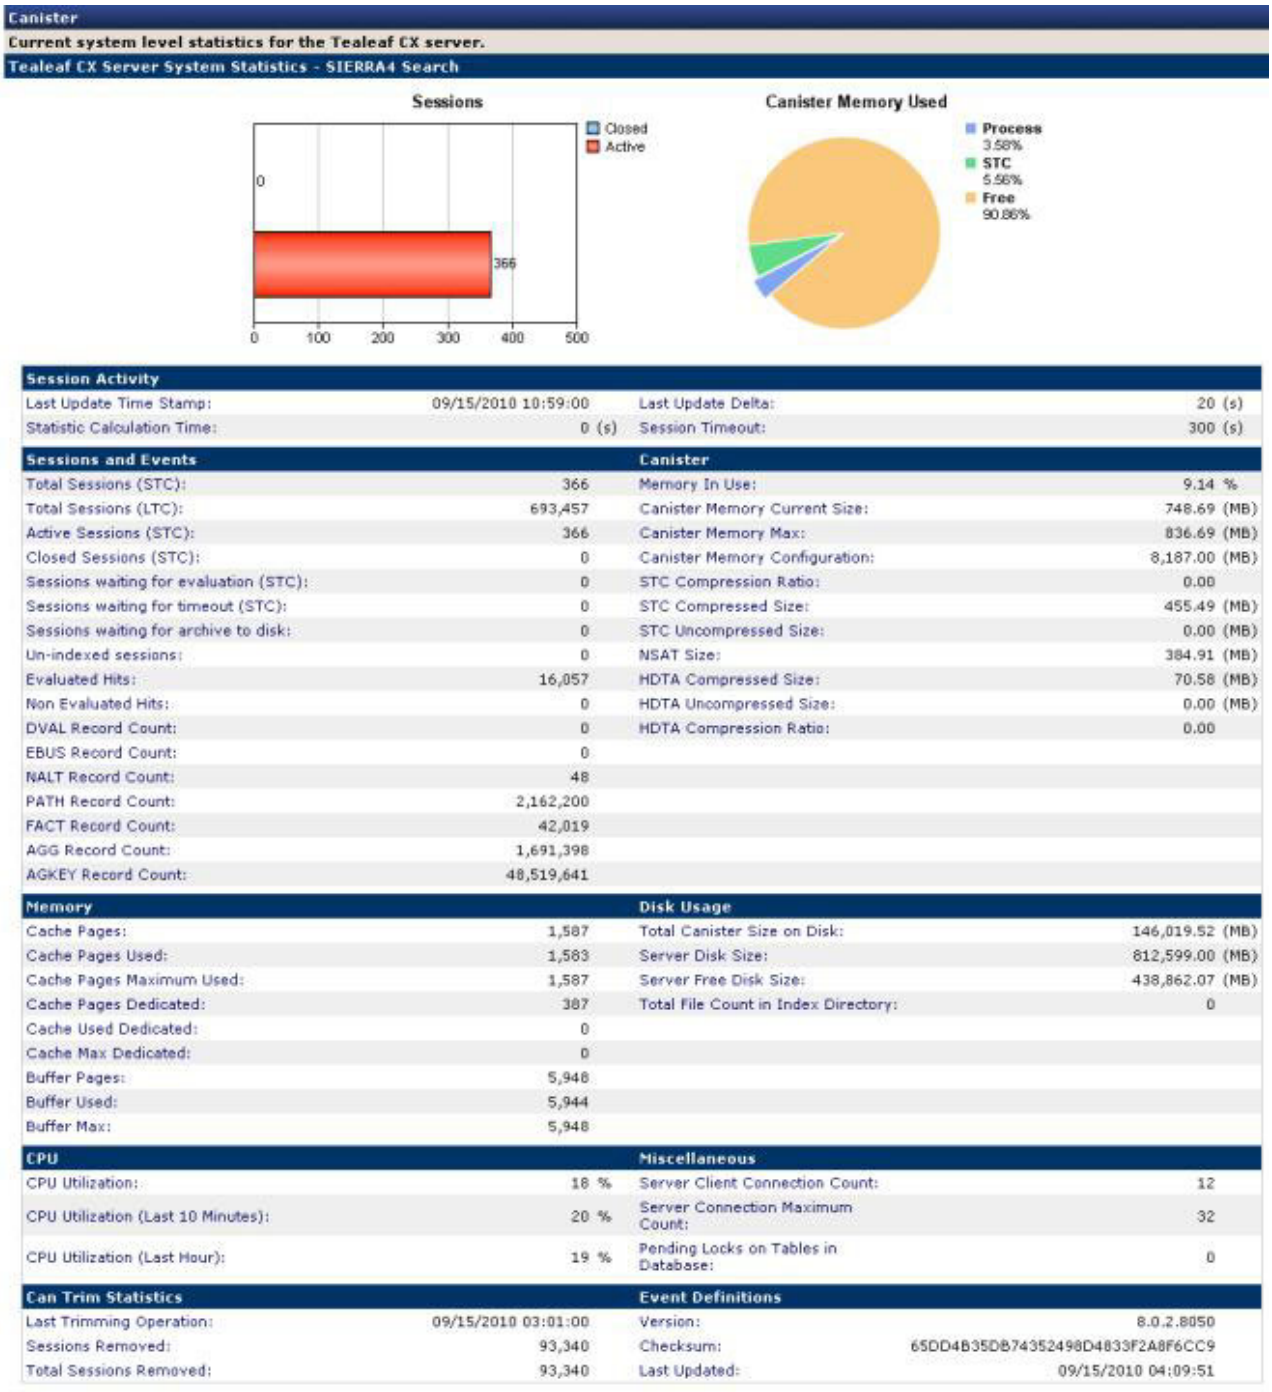

Session Activity			
Last Update Time Stamp:	09/15/2010 10:59:00	Last Update Delta:	20 (s)
Statistic Calculation Time:	0 (s)	Session Timeout:	300 (s)

Sessions and Events		Canister	
Total Sessions (STC):	366	Memory In Use:	9.14 %
Total Sessions (LTC):	693,457	Canister Memory Current Size:	748.69 (MB)
Active Sessions (STC):	366	Canister Memory Max:	836.69 (MB)
Closed Sessions (STC):	0	Canister Memory Configuration:	8,187.00 (MB)
Sessions waiting for evaluation (STC):	0	STC Compression Ratio:	0.00
Sessions waiting for timeout (STC):	0	STC Compressed Size:	455.49 (MB)
Sessions waiting for archive to disk:	0	STC Uncompressed Size:	0.00 (MB)
Un-indexed sessions:	0	NSAT Size:	384.91 (MB)
Evaluated Hits:	16,057	HDTA Compressed Size:	70.58 (MB)
Non Evaluated Hits:	0	HDTA Uncompressed Size:	0.00 (MB)
DVAL Record Count:	0	HDTA Compression Ratio:	0.00
EBUS Record Count:	0		
NALT Record Count:	48		
PATH Record Count:	2,162,200		
FACT Record Count:	42,019		
AGG Record Count:	1,691,398		
AGKEY Record Count:	48,519,641		

Memory		Disk Usage	
Cache Pages:	1,587	Total Canister Size on Disk:	146,019.52 (MB)
Cache Pages Used:	1,583	Server Disk Size:	812,599.00 (MB)
Cache Pages Maximum Used:	1,587	Server Free Disk Size:	438,862.07 (MB)
Cache Pages Dedicated:	387	Total File Count in Index Directory:	0
Cache Used Dedicated:	0		
Cache Max Dedicated:	0		
Buffer Pages:	5,948		
Buffer Used:	5,944		
Buffer Max:	5,948		

CPU		Miscellaneous	
CPU Utilization:	18 %	Server Client Connection Count:	12
CPU Utilization (Last 10 Minutes):	20 %	Server Connection Maximum Count:	32
CPU Utilization (Last Hour):	19 %	Pending Locks on Tables in Database:	0

Can Trim Statistics		Event Definitions	
Last Trimming Operation:	09/15/2010 03:01:00	Version:	8.0.2.8050
Sessions Removed:	93,340	Checksum:	65DD4B35DB74352498D4833F2A8F6CC9
Total Sessions Removed:	93,340	Last Updated:	09/15/2010 04:09:51

Figure 10. Canister Status report - Single Server

Some key fields are described below.

Note: You can create alerts to trigger based on the values for some of these metrics. See "TEM Alerts Tab" in the *IBM Tealeaf Event Manager Manual*.

Session Activity

Statistics on session-related activities.

Field Description

Last Update Time Stamp
Last time that the Canister was queried for statistical data

Statistic Calculation Time
Time in seconds required to compute Canister statistics

Last Update Delta
Time in seconds between last update and delivery of statistics to the Portal.

Session Timeout
Currently configured session timeout setting in seconds for the Canister.

- Session timeout can be overridden by event for individual settings. See "TEM Events Tab" in the *IBM Tealeaf Event Manager Manual*.
- For more information on setting the default session timeout, see "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

Sessions and Events

Counts of sessions and events in the Canister.

Field Description

Total Sessions (STC)
Number of sessions in the Short Term Canister at last update.

Total Sessions (LTC)
Number of sessions in the Long Term Canister at last update.

Active Sessions (STC)
Number of active sessions in the Short Term Canister at last update.

Closed Sessions (STC)
Number of closed sessions in the Short Term Canister at last update.

Sessions waiting evaluation (STC)
Number of sessions waiting for evaluation in the Short Term Canister.

Sessions waiting for timeout (STC)
Number of sessions in the Long Term Canister that have not been updated recently and are waiting for the `Session Timeout` to be reached, when they will be closed.

Sessions waiting for archive to disk
Number of sessions waiting to be written to disk for indexing.

Un-indexed sessions
Number of saved sessions that have not yet been indexed.

Evaluated Hits
Number of hits in the STC that have been evaluated since prior update.

Non Evaluated Hits
Number of hits in the STC that have been not been evaluated since prior update.

DVAL Record Count
Count of records in the DVAL database table, which stores dimension values.

EBUS Record Count
Count of records in the EBUS database table, which stores events queued for the Tealeaf Event Bus.

- The Event Bus is a component of IBM Tealeaf cxConnect for Data Analysis, a separately licensable product of the IBM Tealeaf CX platform. See "Tealeaf Event Bus" in the *IBM Tealeaf cxConnect for Data Analysis Administration Manual*.

NALT Record Count
Count of records in the NALT database table, which stores alert data records.

PATH Record Count
Count of records in the PATH database table, which stores captured path information.

FACT Record Count
Count of records in the FACT database table, which stores fact instance data (event + event value + dimension value).

AGG Record Count
Count of records in the AGG database table, which stores aggregation data.

AGKEY Record Count
Count of records in the AGKEY database table, which stores aggregation keys.

Canister

Data storage information in the Canister.

Field Description

Memory in Use
Current® percentage of available memory in use in the Canister

Canister Memory Current Size
Current size in MB of Canister memory in use

Canister Memory Max
Maximum available Canister memory in MB

Canister Memory Configuration
Maximum available memory in MB on the Canister server

STC Compression Ratio
Compression ratio of memory in the STC

STC Compressed Size
Size of STC memory in MB when compressed

STC Uncompressed Size
Size of STC memory in MB when uncompressed

NSAT Size
Size of session attribute table

HDTA Compressed Size
Compressed size of the hit table

HDTA Uncompressed Size
Uncompressed size of the hit table

HDTA Compression Ratio
Compression ratio of the hit table

Memory

Memory usage in the Canister.

Field Description

Cache Pages
Number of data cache pages configured for use

Cache Pages Used
Number of data cache pages in use

Cache Pages Maximum Used
High-water mark for data cache pages used

Cache Pages Dedicated
Number of data cache pages dedicated to configured tables

Cache Used Dedicated
Number of dedicated pages in use

Cache Max Dedicated
High-water mark for dedicated pages used

Buffer Pages
Number of pages configured for table indexes

Buffer Used
Number of table index pages in use

Buffer Max
High-water mark for index pages in used

Disk Usage

Disk drive usage for the Canister.

Field Description

Total Canister Size on Disk
Total size in MB of Canister data, including indexes, stored in the disk

Server Disk Size
Total size of the hard disk in MB for the server

Server Free Disk Size
Total size in MB of available disk space

Total File Count in Index Directory
This figure is non-zero when the number of indexing workfiles has increased due to an indexing error condition. Typically, this metric is `0`.

CPU

Current and recent CPU usage for the Canister.

Field **Description**

CPU Utilization
Percentage of CPU resources in use at last update

CPU Utilization (Last 10 Minutes)
Percentage of CPU resources in use in 10 minutes prior to last update

CPU Utilization (Last Hour)
Percentage of CPU resources in use in last hour prior to last update

Miscellaneous

Miscellaneous server information.

Field **Description**

Server Client Connection Count
Current count of connections to the Canister server

Server Connection Maximum Count
Maximum count of simultaneous connections permitted to the Canister server

Pending Locks on Tables in Database
Count of lock requests pending for tables in the LSSN database

Can Trim Statistics

Statistics on the last Canister trimming operation.

- If CanTrim has never been executed on the Canister, these values are listed as N/A.

Field **Description**

Last Trimming Operation
Timestamp of last trimming operation of Canister data

Sessions Removed
Number of sessions trimmed in last CanTrim operation

Total Sessions Removed
Total number of sessions trimmed since last server restart

Event Definitions

Version information for the event definitions currently in use by the Canister.

- This information is provided for debugging purposes only.

Field **Description**

Version
Tealeaf build number for the event definitions

Checksum
A checksum value for the event definitions file on the server

Last Updated
Timestamp for when the event definitions were last updated in the Canister

DecoupleEx

This report displays the operational status of the Tealeaf Transport Service and Short Term Canister in the IBM Tealeaf CX datastore. You can use the statistics on this page as a guide for troubleshooting problems with the Short Term Canister.

- For more information on configuring the Transport Service, see Chapter 4, “Tealeaf Management System,” on page 259.
- For more information on tuning the Short Term Canister, see "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.
- For more information on the DecoupleEx session agent, see "Extended Decoupler Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

All servers

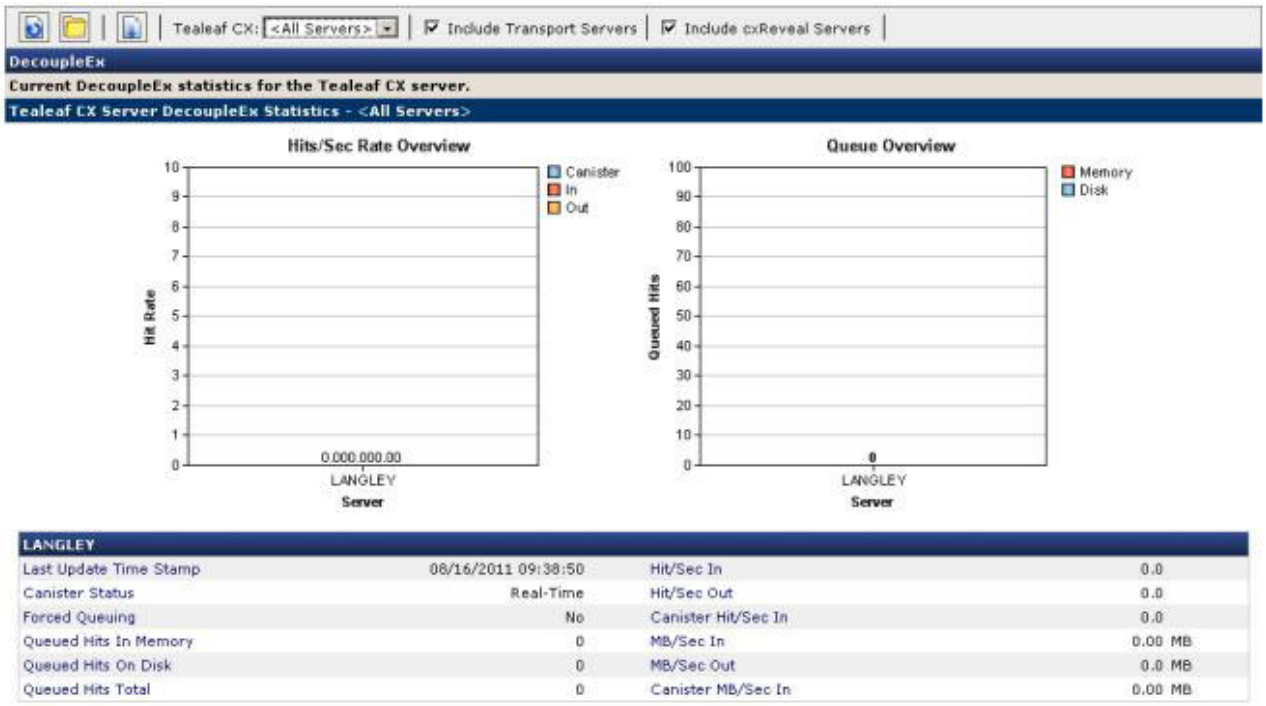

Figure 11. DecoupleEx Status report - All Servers

- To display status information from the Transport Servers in the report, click the Include Transport Servers checkbox.
- To display status information from the IBM Tealeaf cxReveal Servers in the report, click the Include IBM Tealeaf cxReveal Servers checkbox.

  Note: IBM Tealeaf cxReveal Servers are used to capture session attribute information and store it in a separate database for immediate and rapid session search. IBM Tealeaf cxReveal is a separately licensable product of the Tealeaf platform and requires a separate installation. See "cxReveal Installation" in the *IBM Tealeaf cxReveal Administration Manual*.

Field Description

`Last Update Time Stamp`
A timestamp indicating the last time this report was updated.

`Canister Status`
Current status of the canister: `Real-Time` or `Spooling`.

Forced Queuing
Queuing of hits is enforced automatically. When set to `yes`, a start queue command was issued, and hits are being queued even though the Short Term Canister has not crossed any thresholds.

Queued Hits in Memory
The number of hits queued in memory.

Queued Hits on Disk
The number of hits queued on disk.

Queued Hits Total
The total number of queued hits.

Hits/Sec In
The number of hits being submitted to DecouplerEx per second. This value shows the rate of hits per second as they enter the extended decoupler providing you with constant snapshots of session data as it travels through the Transport pipeline.

- If the rate of hits entering the Extended Decoupler equals the rate coming out, then the Transport Pipeline is in good health and fully operational.
- If there are more hits coming in than coming out, something may be wrong in the pipeline or Short Term Canister, or queuing may have started.

Hits/Sec Out
This value indicates the number of hits per second exiting the Extended Decoupler. These values fluctuate and may not always match incoming hits, but they should remain relatively consistent with incoming hit values.

Canister Hits/Sec In
The number of hits being submitted to the downstream canister per second.

MB/Sec In
The megabytes of data submitted per second. This value essentially measures the health of the Transport pipeline. The megabytes of data coming in should generally match the megabytes of data coming out of the Extended Decoupler. If they do not, there is a problem either in the Transport Pipeline or in the Short Term Canister.

MB/Sec Out
The megabytes of data outputted per second exiting the Extended Decoupler. These values fluctuate and may not always match incoming hits, but they should remain relatively consistent with incoming hit values.

Canister MB/Sec In
The megabytes per second of data forwarded to the canister.

Single server

When a single DecoupleEx server is selected, the following report is displayed.

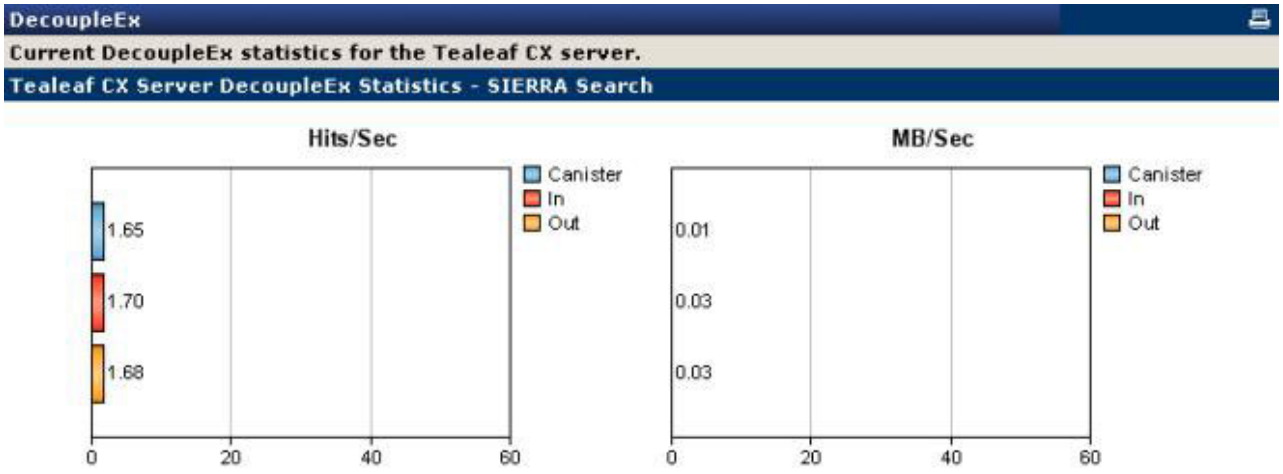

Status			
DecoupleEx Status:	Memory	Last Update Time Stamp:	03/31/2010 12:28:10
Canister Status:	Real-Time	Canister Message:	Canister is operational
Time of Hits on Disk:	0	Total Dropped Hits:	0
Forced Queuing:	No	Current Spooled Hits:	1
Canister Hits Ratio:	0.98	Hits In/Out Ratio:	1.01
Canister Hits In		**Canister Megabytes In**	
Hits/Sec In Current:	1.6	MB/Sec In Current:	0.01 (MB)
Hits/Sec In Total:	1.6	MB/Sec In Total:	0.01 (MB)
Hits/Sec In Max:	2.4	MB/Sec In Max:	0.05 (MB)
Hits/Sec In Min:	1.2	MB/Sec In Min:	0.00 (MB)
Hits In		**MegaBytes In**	
Hits/Sec In Current:	1.7	MB/Sec In Current:	0.03 (MB)
Hits/Sec In Total:	1.6	MB/Sec In Total:	0.03 (MB)
Hits/Sec In Max:	2.5	MB/Sec In Max:	0.06 (MB)
Hits/Sec In Min:	1.2	MB/Sec In Min:	0.01 (MB)
Hits Out		**MegaBytes Out**	
Hits/Sec Out Current:	1.7	MB/Sec Out Current:	0.03 (MB)
Hits/Sec Out Total:	1.6	MB/Sec Out Total:	0.03 (MB)
Hits/Sec Out Max:	2.5	MB/Sec Out Max:	0.06 (MB)
Hits/Sec Out Min:	1.2	MB/Sec Out Min:	0.01 (MB)
Total Hits/MegaBytes		**Current Sample Hits/MegaBytes**	
Canister In Hits:	114,537	Canister In Hits:	99
Canister In MB:	534.06 (MB)	Canister In MB:	0.48 (MB)
In Hits:	116,501	In Hits:	102
In MB:	1,830.40 (MB)	In MB:	1.80 (MB)
Out Hits:	116,552	Out Hits:	101
Out MB:	1,831.29 (MB)	Out MB:	1.79 (MB)
Current Memory/Disk		**Max Memory/Disk**	
Hits In Memory:	1	Hits In Memory:	52
Hits On Disk:	0	Hits On Disk:	52
Hits Total:	1	Hits Total:	53
MB In Memory:	0.02 (MB)	MB In Memory:	1.34 (MB)
MB On Disk:	0.00 (MB)	MB On Disk:	0.90 (MB)
MB Total:	0.02 (MB)	MB Total:	1.34 (MB)
Time of First Hit in Spool Files:	N/A		

Figure 12. DecoupleEx Status report - Single Server

Some key fields are described below.

Status

Field Description

DecoupleEx Status
The current status of the DecoupleEx session agent:

- Memory - hits are being processed in memory only.
- Disk - hits are being spooled to disk because memory is full.
- UNKNOWN - the status of the DecoupleEx session agent is not known. The session agent is not operational.

Canister Status
Current status of the canister: UP or DOWN.

Time of Hits On Disk
The average time that hits are spending on disk.

Forced Queuing
Queuing of hits is enforced automatically. When set to yes, a start queue command was issued, and hits are being queued even though the Short Term Canister has not crossed any thresholds.

Canister Hits Ratio
The ratio of hits entering the canister and exiting the canister.

Last Update Time Stamp
A timestamp indicating the last time this report was updated.

Canister Message
The current status message from the canister.

Total Dropped Hits
The total number of hits that have been dropped by the canister since last restart.

Current Spooled Hits
The current number of hits that have been spooled to disk.

Hits In/Hits Out Ratio
The ratio of hits entering the DecoupleEx session agent and exiting the session agent.

Notes® on minimum, maximum and total values

In this report, these values are computed since the last time the Transport Service was restarted.

- Maximum values may be skewed. When the Transport Service is restarted, all hits that have been queued by the PCA are submitted at the same time.

Canister Hits In and Canister MegaBytes In

The current, total, maximum and minimum rates of traffic entering the canister per second, as measured in number of hits and megabytes of data.

Hits In and MegaBytes In

The current, total, maximum and minimum rates of traffic entering the DecoupleEx session agent per second, as measured in number of hits and megabytes of data.

Hits Out and MegaBytes Out

The current, total, maximum and minimum rates of traffic exiting the DecoupleEx session agent per second, as measured in number of hits and megabytes of data.

Total Hits/MegaBytes and Current Sample Hits/MegaBytes

The total number of hits and megabytes of data entering the canister, entering the DecoupleEx session agent, and exiting the session agent.

- The period for the current sample is approximately one minute.

Current Memory/Disk and Max Memory/Disk

The current number of hits in memory, queued on disk, and total in both areas, as well as the megabytes of storage they occupy.

The same metrics are listed for the maximum values over the reporting period.

`Time of First Hit in Spool Files` indicates the timestamp when the first hit was spooled, which indicates an overflow condition, since the last restart of the pipeline.

- If this value is set to `N/A`, no hits have been spooled since the last restart.

Events

Note: This report has been deprecated as of Release 8.0. See "Event Activity" in the *IBM Tealeaf Reporting Guide*.

Storage

Displays information about the Long Term Archive files and indexes for the Canister by server. These charts provide a quick glance at the system usage.

- The per server graph displays the last 7 days of session count and index/archive size ratios.
- Indexes should have a ratio < 1.

Common data

The following fields are common to the All Server and Single Server Storage reports:

Field Description

`Server`
The name of the Tealeaf server.

`Index Style`
The method by which data is indexed.

`Total Days`
The total number of days of data that are currently indexed.

`LSSN Space Used (MB)`
The amount of space used by the Long Term Canister in megabytes.

`LSSN Space Free (MB)`
The amount of free space available to the Long Term Canister in megabytes.

LSSN Disk Size (MB)
The size of the disk in megabytes where LSSN files reside.

Index Space Used (MB)
The amount of space used by the index in megabytes.

Index Space Free (MB)
The amount of free space available to the index in megabytes.

Index Disk Size (MB)
The amount of space used on the disk by the index in megabytes.

Total Space Used (MB)
The total space used by the Long Term Canister and the indexes in megabytes.

All Servers

Storage
Storage statistics for the Tealeaf CX server.
Tealeaf CX Server Storage Statistics - <All Servers>
Server | Index Style | Total Days | LSSN Space Used (MB) | LSSN Space Free (MB) | LSSN Disk Size (MB) | Index Space Used (MB) | Index Space Free (MB) | Index Disk Size (MB) | Total Space Used (MB)
SIERRA4 Search | FastRsp | 7 | 84,209.3 (10.36%) | 383,963.5 (47.25%) | 812,599.0 | 22,542.9 (2.77%) | 383,963.5 (47.25%) | 812,599.0 | 106,752.2
Total | | | 84,209.3 | | | 22,542.9 | | | 106,752.2
Avg. Total Index/LSSN Size (MB) Per Day
Index Size
LSSN Size
18,000
16,000
14,000
12,000
10,000
8,000
6,000
4,000
2,000
0
Size (MB)
15,250.3
3,220.4
12,029.9
SIERRA4 Search
Server
Server | Avg. LSSN Size (MB) Per Day | Avg. Index Size (MB) Per Day | Total Avg. Size (MB) Per Day
SIERRA4 Search | 12,029.9 | 3,220.4 | 15,250.3

Figure 13. Storage Status report - All Servers

Field Description

Server
The name of the Tealeaf server.

Avg. LSSN Size (MB) Per Day
Average size in megabytes of the LSSN file or files per day.

Avg. Index Size (MB) Per Day
Average size in megabytes of the LSSN index file or files per day.

Total Avg. Size (MB) Per Day
Sum of the other two columns.

Single Server

In the details for an individual server, you can review storage data from the selected server for each day over a preceding period.

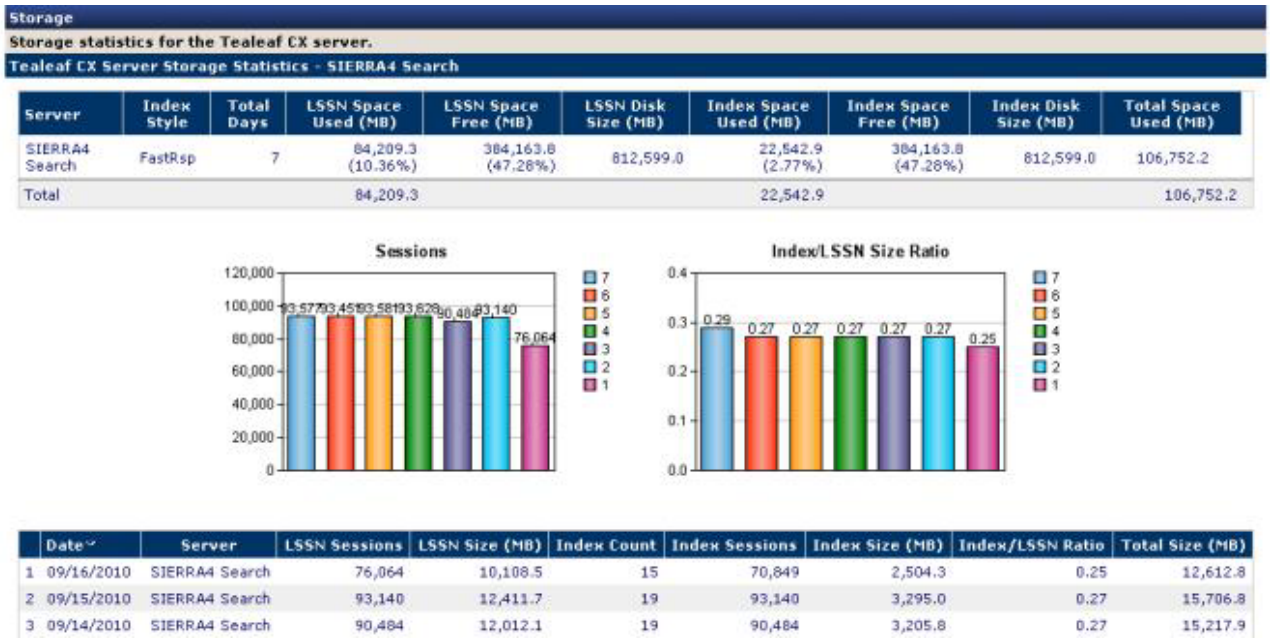
Storage
Storage statistics for the Tealeaf CX server.
Tealeaf CX Server Storage Statistics - SIERRA4 Search
Server | Index Style | Total Days | LSSN Space Used (MB) | LSSN Space Free (MB) | LSSN Disk Size (MB) | Index Space Used (MB) | Index Space Free (MB) | Index Disk Size (MB) | Total Space Used (MB)
SIERRA4 Search | FastRsp | 7 | 84,209.3 (10.36%) | 384,163.8 (47.28%) | 812,599.0 | 22,542.9 (2.77%) | 384,163.8 (47.28%) | 812,599.0 | 106,752.2
Total | | | 84,209.3 | | | 22,542.9 | | | 106,752.2
Sessions
120,000
100,000
80,000
60,000
40,000
20,000
0
93,577
93,140
76,064
7
6
5
4
3
2
1
Index/LSSN Size Ratio
0.4
0.3
0.2
0.1
0.0
0.29
0.27
0.27
0.27
0.27
0.27
0.25
7
6
5
4
3
2
1
Date | Server | LSSN Sessions | LSSN Size (MB) | Index Count | Index Sessions | Index Size (MB) | Index/LSSN Ratio | Total Size (MB)
1 09/16/2010 SIERRA4 Search 76,064 10,108.5 15 70,849 2,504.3 0.25 12,612.8
2 09/15/2010 SIERRA4 Search 93,140 12,411.7 19 93,140 3,295.0 0.27 15,706.8
3 09/14/2010 SIERRA4 Search 90,484 12,012.1 19 90,484 3,205.8 0.27 15,217.9

Figure 14. Storage Status report - Single Server

Some key fields are described below.

Field **Description**

Date Date to which the row data applies.

- To review information on the indexes used on the selected date, click the drill-down link. See "Indexes" on page 29.

Server
Name of the server.

LSSN Sessions
Number of sessions stored in the LSSN for that day.

LSSN Size (MB)
Size of the LSSN in megabytes for that day.

Index Count
Number of indexes for the LSSN.

Index Sessions
Number of sessions indexed in all of the indexes.

Index Size (MB)
Total size of the indexes for the date.

Index/LSSN Ratio
The ratio of the size of the index compared to the size of the LSSN, which indicates the efficiency of the indexes.

Total Size (MB)
The total storage size in megabytes for the date. This value equals the sum of LSSN Size (MB) and Index Size (MB).

Indexes

When an individual date is clicked in the Single Server storage report, the following popup is displayed:

Figure 15. System Status - Storage Indexes

Field	Description
Day	Date of the index report
Julian	Date in the Julian calendar of the index report
First Use	Timestamp of first usage on the selected date
Last Use	Timestamp of last usage on the selected date

Below the report header, you can read report data on individual indexes on the server.

Field	Description
Path	Directory where the index is stored
Valid	If True, the index passed its last index check.
In Use	If True, the Tealeaf indexer is adding sessions to the index.
Index Size	The size in megabytes of the index on the disk
First Use	Timestamp of first usage of the index on the selected date
Last Use	Timestamp of last usage of the index on the selected date
First Session	Session ID of the first session stored in the index
Last Session	Session ID of the last session stored in the index
Check Required	If True, the index must pass an index check before it can be used.
Doc Count	Number of documents in the index

Health Based Routing (HBR)

This report displays the current status of the canisters monitored by Health Based Routing.

- **Health Based Routing** enables the routing of hit data to multiple canisters for processing based on the health and load on each canister. For more information on HBR, see "Health-Based Routing (HBR) Session Agent" in the *IBM Tealeaf CX Configuration Manual*.
- Through the Portal Management page, you can add and configure Transport Servers. For more information on how to configure a transport server, see “Managing Tealeaf Servers” on page 115.

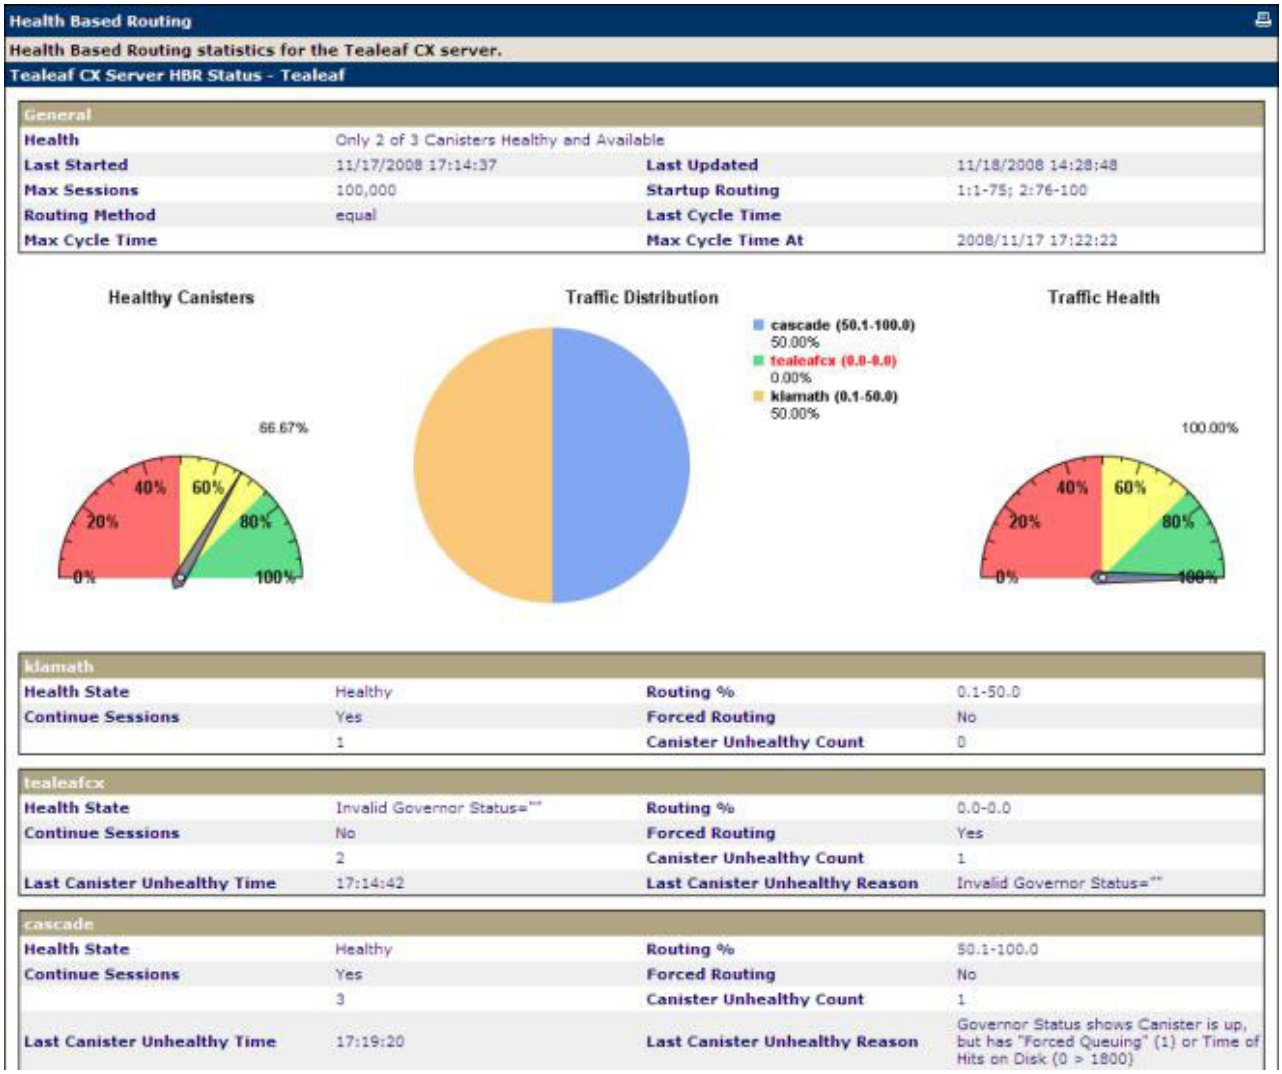

Figure 16. Health Based Routing (HBR) Status report

General

Field **Description**

`Health`
A message indicating the current state of all canisters. A properly functioning HBR system is indicated by the message `Healthy`.

`Last Started`
A timestamp indicating the last time HBR was started.

`Max Sessions`
The maximum number of sessions that HBR will manage at any one time. HBR has an internal session table to indicate which canister a session belongs and this value determines the size of that table. Entries are managed via a LRU (Least Recently Used) algorithm, i.e., if the table is full and a new entry is needed the oldest entry is thrown out. For most sites the default value of 100,000 is sufficient.

`Routing Method`
The method by which sessions are routed to individual canisters. The most common value of "equal" is somewhat of a misnomer as canisters can be

rated differently ("equal" being a holdover from early HBR versions which required all target canisters were equal).

Max Cycle Time
The longest time required to poll all the target canisters for their health. Normally this value will be in the single digit second value.

Current Spooled Hits
Number of hits that are currently spooled by the HBR server.

Last Updated
Timestamp for when HBR last updated this information.

Startup Routing
A fallback (safety valve) value if for some reason HBR can not determine the health status of the target canisters. This should never come into play as if HBR can not determine the health of a canister it will assume it to be unhealthy and keep it out of rotation.

Last Cycle Time
A timestamp indicating the last time that the canisters were restarted.

Max Cycle Time At
A timestamp indicating when the `Max Cycle Time` value occurred.

Time of First Hit in Spool Files
A timestamp indicating when the first hit was added to the currently open spool files. This value is maintained until the spool files are removed from the disk.

Individual Canisters

Some key fields are described below.

Field Description

Health State
A message indicating the current state of the canister. A properly functioning canister is indicated by the message `Healthy`.

Continue Sessions
If set to `yes`, the canister attempts to merge fragmented sessions into whole sessions.

Governor Ex Status Instance
An identifier internal to HBR of each instance of the DecoupleEx session agent in the Tealeaf system.

Routing %
If HBR is configured to route traffic based on percentages, this value indicates the percentage range of traffic that is routed to this canister.

Forced Routing
If set to `Yes`, some web traffic is forced to be processed by this canister.

Canister Unhealthy Count
Over the previous reporting period, this value indicates the number of times that the canister has been in an unhealthy state.

Last Canister Unhealthy Time
A timestamp indicating the last time that the canister was in a non-`Healthy` state.

Last Canister Unhealthy Reason
A readable message identifying the reason why the canister was unhealthy at Last Canister Unhealthy Time.

Database Filegroup Size

The Database Filegroup Size report contains useful metrics on the current size and growth of the Tealeaf databases.

- The data collector polls for this data once per day at 2:00 AM.
- See "CX Databases" in the *IBM Tealeaf Databases Guide*.
- See "cxResults Databases" in the *IBM Tealeaf Databases Guide*.
- For more information on troubleshooting databases, see "Troubleshooting - Tealeaf Databases" in the *IBM Tealeaf Troubleshooting Guide*.
  - If this report does not contain information on free space or unused space, there may be a permissions issue in SQL Server. See "Troubleshooting - Tealeaf Databases" in the *IBM Tealeaf Troubleshooting Guide*.

Note: After installation or upgrade is completed, data may not be available in this report until up to 1 hour later. If data does not appear in the report after 1 hour, there is a known issue that may be affecting the updating of the report. See "Troubleshooting - Reporting" in the *IBM Tealeaf Troubleshooting Guide*.

Note: To be able to retrieve all of the available data to populate this report, the TLADMIN account requires additional permissions. See "SQL Server Authentication" in the *IBM Tealeaf Databases Guide*.

Data displayed in this report is specific to the currently selected database on the selected server for the specified date range.

Figure 17. Database Filegroup Size

In the chart, you can review the current size in megabytes of each filegroup in the currently selected database.

- To review size information for another Tealeaf database, make your selection from the database drop-down. You must select the server where the database is hosted.

  Note: Since some databases, such as the TL_SEARCH database, may have multiple instances on separate servers in your Tealeaf environment, please verify that you have selected the appropriate server.

- The list of available servers includes all active and inactive Report Servers and IBM Tealeaf cxReveal Database Search Servers, as defined in the Portal Management page. See “Managing Tealeaf Servers” on page 115.

If the date range is more than a single day, the Daily Percentage Growth chart is added to the right of the Filegroup Size chart. Charted growth for the selected database is displayed across the configured date range.

Available databases include the following:

Note: The following database names are the defaults provided by Tealeaf. The names may differ in your Tealeaf solution.

Default DB Name
Description

TL_SYSTEM
The System database is used for configuring and operating the Portal application and Tealeaf reporting capabilities. See "CX Databases" in the *IBM Tealeaf Databases Guide*.

TL_REPORTS
The Reporting database contains statistics collected from Tealeaf Processing Servers, including event, path, and statistics information in aggregated format. See "CX Databases" in the *IBM Tealeaf Databases Guide*.

TL_RSEXTRACTOR
The RS Extractor database contains the results of segment analysis. See "CX Databases" in the *IBM Tealeaf Databases Guide*.

TL_STATISTICS
The Statistics database contains performance information on the Tealeaf system. See "CX Databases" in the *IBM Tealeaf Databases Guide*.

TL_VISREPORT
The Visitor Reporting database contains visitor data stored for reporting purposes. See "cxResults Databases" in the *IBM Tealeaf Databases Guide*.

TL_VISSTAGE
The Visitor Staging database is used to contain data before it is inserted into the Visitor Reporting database. See "cxResults Databases" in the *IBM Tealeaf Databases Guide*.

TL_SEARCH
The Search database is used by IBM Tealeaf cxReveal to maintain session attribute information for search purposes.

- To report information from all IBM Tealeaf cxReveal servers, select `<All Servers>`. Optionally, you can select an individual server from the drop-down.

  Note: This report is not populated if IBM Tealeaf cxReveal session attribute search and the IBM Tealeaf cxReveal search database have not been installed. See "Configuring Session Attribute Search" in the *IBM Tealeaf cxReveal Administration Manual*.

  See "cxReveal Database" in the *IBM Tealeaf Databases Guide*.

In the detail table, you can review statistics on each filegroup of the currently selected database for each date in the selected range.

Field **Description**

Date Date when report information was collected for the filegroup

File Group
The filegroup to which the row of data applies

Total Space (MB)
Total space occupied by the filegroup on the storage device on the listed date

Free Space (MB)
Total space in megabytes on the storage device

% Space Free
Percentage of space available on the storage device

Size (MB)
Total space in megabytes occupied by the filegroup

Daily % Growth
Percentage of growth in database size from the previous day

- This value is zero if the date range is a single date or if no data is available.

Avg Daily Growth
Average daily growth of database size in megabytes over the selected date range

Daily Delta
Change in database size in megabytes from the previous day

Monthly % Growth
Percentage of growth in database size from the previous month, if the date range is at least two months

Avg Monthly Growth
Average monthly growth of database size in megabytes if the date range is at least two months

Monthly Delta
Change in database size in megabytes from the previous month

Location
Filename of the database file

Database Table Size

This report provides statistical information on the current size and growth of database tables based on the type of data in them.

- See "CX Databases" in the *IBM Tealeaf Databases Guide*.
- See "cxResults Databases" in the *IBM Tealeaf Databases Guide*.
- For more information on troubleshooting databases, see "Troubleshooting - Tealeaf Databases" in the *IBM Tealeaf Troubleshooting Guide*.
  - If this report does not contain information on free space or unused space, there may be a permissions issue in SQL Server. See "Troubleshooting - Tealeaf Databases" in the *IBM Tealeaf Troubleshooting Guide*.

The data collector polls for this data at once a day at 2:00 AM.

Note: After installation or upgrade is completed, data may not be available in this report until up to 1 hour later.

- Data displayed in this report is specific to the currently selected data group and date range.

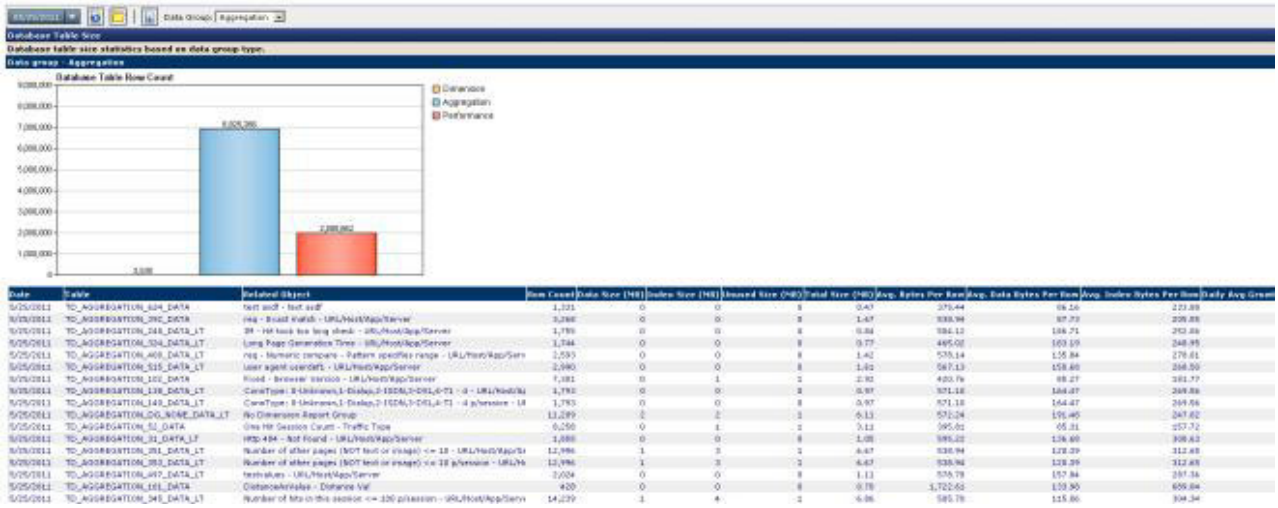

Figure 18. Database Table Size

If the date range is more than a single day, the Daily Percentage Growth chart is added to the right of the Row Count chart. Charted growth for the selected data group is displayed across the configured date range.

A data group is a loose organization of database tables based upon the type of data that they contain. Available data groups include the following:

Data Group
Description

`Aggregation`
Data that has been collected by the Data Collector from each Canister in the environment and then aggregated for reporting purposes.

`Dimension`
Fact data (event + dimension value) stored for each occurrence of each event.

`Performance`
Performance data is stored for individual hits captured by the Tealeaf system

For the selected data group, the detail table is updated with statistics from each relevant database table for each date in the selected date range. Listed fields include the following:

Field **Description**

`Date` Date when report information was collected for the database table

`Table` The database table to which the row of data applies

`Related Object`
The system object that is associated with the table, depending on the data group:

- For Dimension tables: the Dimension name.
- For Aggregation tables: event and maybe report group name.

- For Performance tables: no related objects

Row Count
Number of rows in the table for the listed date

Data Size (MB)
Size of data in megabytes in the table for the listed date

Index Size (MB)
Size of index data in megabytes in the table for the listed date

Unused Size (MB)
Size available for the table in megabytes for the listed date

Total Size (MB)
Total size of database table in megabytes, including data and index for the listed date

Avg Bytes per Row
Average number of bytes per table row

Avg Data Bytes per Row
Average number of data bytes per table row

Avg Index Bytes per Row
Average number of index bytes per table row

Daily Avg Growth (MB)
Average daily growth of database table in megabytes over the selected date range

Daily % Growth
Percentage of growth in database table size from the previous day

- This value is zero if the date range is a single date or if no data is available.

Daily Delta (MB)
Change in database table size in megabytes from the previous day

Daily Delta Rows
Change in number of rows in the database table from the previous day

Monthly % Growth
Percentage of growth in database table size from the previous month

- This value is zero if the date range is less than two months or if no data is available.

Monthly Delta (MB)
Change in database table size in megabytes from the previous month

Monthly Delta Rows
Change in number of rows in the database table from the previous month

Collection Status

The Collection Status report displays data collection information about the currently configured Tealeaf servers and canisters. The report provides an overview of how fresh is the data in the system and indicates whether report and session data is falling behind.

- To review the Collection Status report, select **Tealeaf** > **System Status** > **Collection Status**.

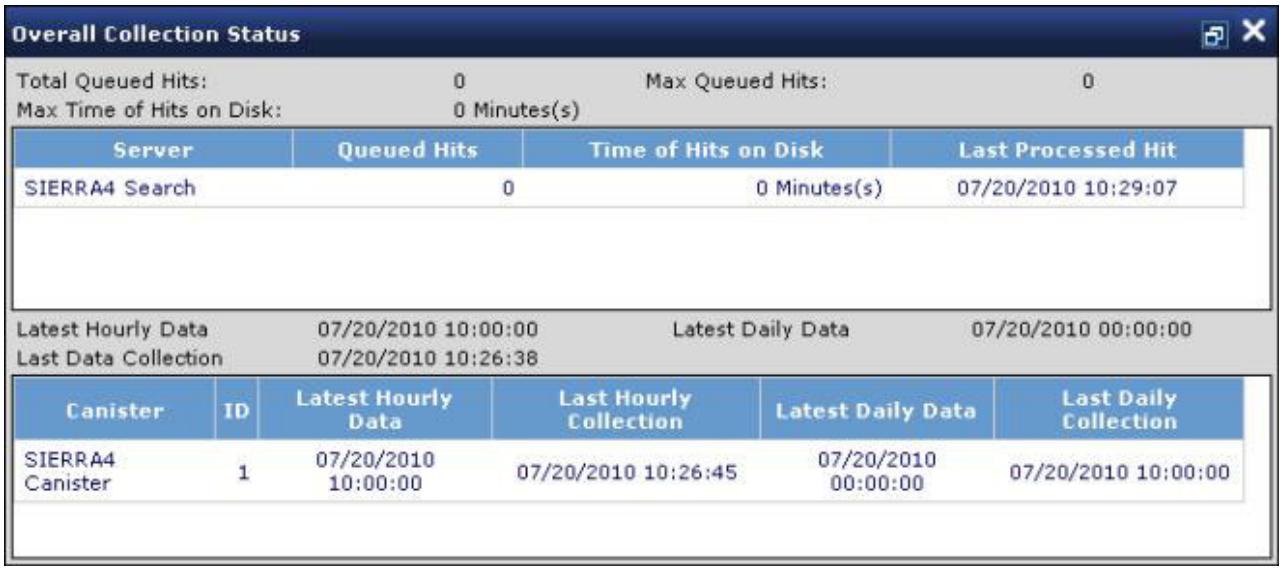

Figure 19. Collection Status

Servers

Field **Description**

Total Queued Hits
The total number of queued hits on all servers.

Max Time of Hits on Disk
The longest period of time a queued hit has remained on disk.

Max Queued Hits
The highest number of hits that were queued on disk over the reporting period.

Server
The name of the individual server.

Queued Hits
The number of hits queued on the server.

Time of Hits on Disk
The average time that hits are queued on disk.

Last Processed Hit
A timestamp indicating when the last hit was processed.

Canisters

Field **Description**

Latest Hourly Data
A timestamp indicating the latest hourly data stored in the system.

Last Data Collection
A timestamp indicating when data was last collected from the canister.

Latest Daily Data
A timestamp indicating the latest daily data stored in the system.

Canister
The name of the canister.

ID The internal identifier for the canister.

Latest Hourly Data
A timestamp indicating the latest hourly data stored in the specified canister.

Last Hourly Collection
A timestamp indicating the most recent collection of hourly data from the specified canister.

Latest Daily Data
A timestamp indicating the latest daily data stored in the specified canister.

Last Daily Collection
A timestamp indicating the most recent collection of daily data from the specified canister.

System Statistics

The System Statistics reports expose a variety of metrics on the data that Tealeaf is currently capturing.

Note: Data in these reports is populated by sending statistic hits from the components listed below. For each component, Statistics hits must be enabled and properly configured. See "Enabling System Statistics Hits."

There are three classes of components: DecoupleEx, IBM Tealeaf CX Capture, and CanisterAS machines.

- `DecoupleEx` - monitors data that is passed to the DecoupleEx session agent at the beginning of the main processing pipeline.
  - See "Extended Decoupler Session Agent" in the *IBM Tealeaf CX Configuration Manual*.
- `CX Capture` - monitors data captured by the IBM Tealeaf CX Passive Capture Application
  - See "PCA Web Console - Statistics Tab" in the *IBM Tealeaf Passive Capture Application Manual*.
- `CanisterAS` - Monitors activity status on the Short-Term and Long-Term Canisters.
  - See "Canister Statistics" on page 57.

For each component class, you can review a series of reports in the Statistics section.

Enabling System Statistics Hits

Before you generate system statistics reports, you must enable the delivery of statistics hits from the relevant components through the Windows pipeline for analysis and reporting.

Steps: Use the following steps to enable statistics hits for the listed components.

1. **Capture statistics hits:** IBM Tealeaf CX Passive Capture Application statistics hits can be enabled through the Delivery tab of the PCA Web Console. See "PCA Web Console - Delivery Tab" in the *IBM Tealeaf Passive Capture Application Manual*.
   - In the Delivery tab, you must set the `Enabled` flag and specify the hostname and port number of the Windows pipeline to receive the statistics hits.

2. **Extended Decoupler statistics hits:** Statistics hits from the Extended Decoupler session agent can be enabled through the TMS Pipeline Editor. It must be enabled for each instance of the session agent in each pipeline. See Chapter 8, “TMS Pipeline Editor,” on page 307.
   - In the Extended Decoupler configuration, you must set the `Generate Stats Hits` setting to `ON`. See "Extended Decoupler Session Agent" in the *IBM Tealeaf CX Configuration Manual*.
3. **PCA statistics hits:** Through the PCA Web Console, you can enable and configure the delivery of PCA stats hits to the Tealeaf Transport Service, through which they are passed into the Windows pipeline for processing. See "PCA Web Console - Delivery Tab" in the *IBM Tealeaf Passive Capture Application Manual*.

Statistics Logging

Statistics hits are inserted into the Tealeaf database by the Stats Logger session agent. This session agent must be enabled and configured in your Windows pipeline.

- See "Statistics Logger Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

Generating a System Statistics Report

To generate a system statistics report:

1. In the Parameters pane, specify the time period. The date and time selector is used to filter the report data to a specific time period. By default, reports are generated from the beginning of the current day to the most recent hour.

Figure 20. System Statistics Parameters

2. In the Component pane, underneath each type of component, you can review and select the servers from which to report statistics:

Figure 21. System Statistics Components

   - To expand all nodes, click **Expand**. To collapse them, click **Collapse**.

3. In the Report pane, you can select the system statistics to include in each report:

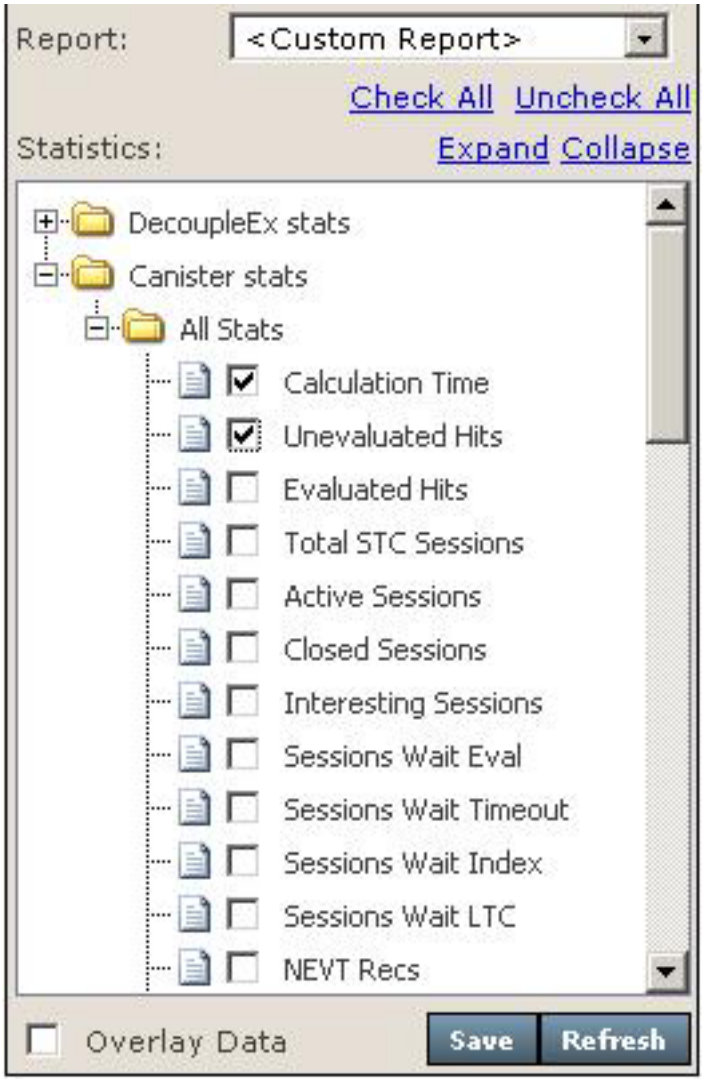

Figure 22. System Statistics Reports

- A series of pre-configured reports are provided through the Reports drop-down. You can include a pre-configured report by selecting it from the drop-down box. See "Pre-Configured Reports" on page 42.
- To expand all nodes, click **Expand**. To collapse them, click **Collapse**.
- To select all reports currently displayed in the pane, click **Check All**. To deselect any selections, click **Uncheck All**.
- To overlay data onto the report, select the Overlay Data checkbox.
- To save the report for later use, click **Save**. See "Saving System Statistics Reports" on page 42.

4. To generate or update the report, click **Refresh**. The updated report is displayed in the right pane.

Saving System Statistics Reports

After you have configured a report, you can save the report configuration for later use.

- To save a report, click **Save** at the bottom of the configuration pane. Enter a name for the report, and click **Save**. The report is saved as a user-defined system statistics report.

User-Defined Reports

To access saved reports at a later time, select the name from the `<User Defined Reports>` category in the Report drop-down. Customized reports are available to all Tealeaf users who can access this screen.

Pre-Configured Reports

This section describes the pre-configured reports.

Component Status

This read-only report provides start and end timestamps for each server component in the Tealeaf environment.

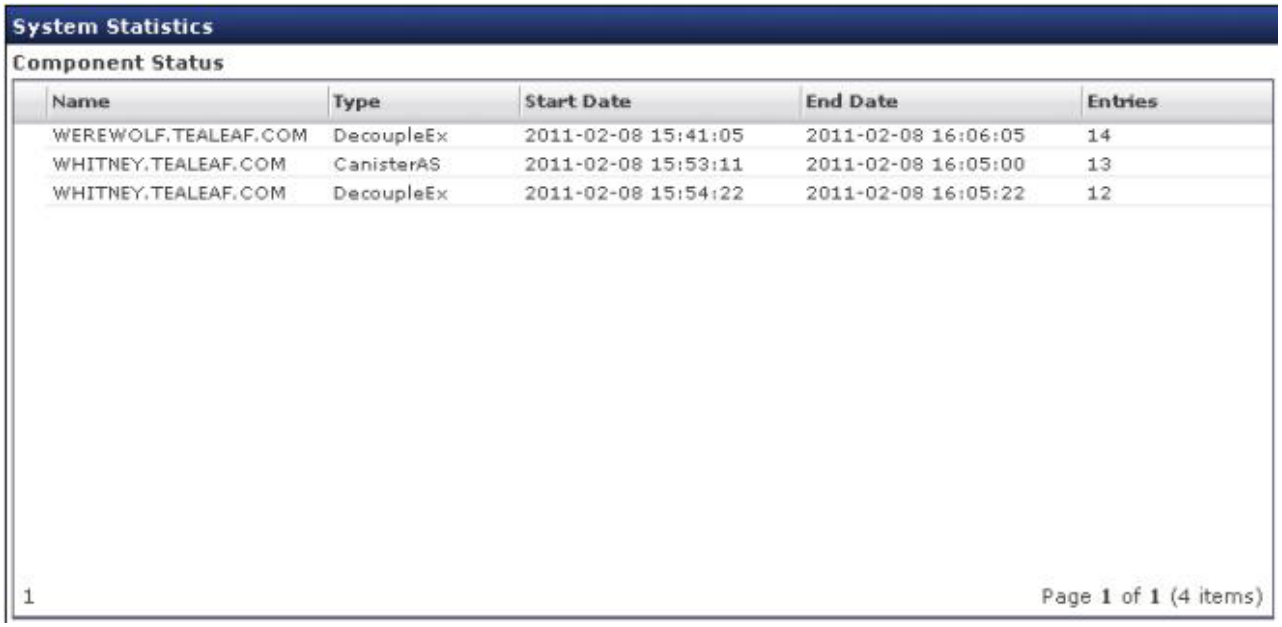

Figure 23. Custom Report - Component Status

Hit Flow

Provides information on the flow of hits in and out of each canister in the system, with statistical summary information at the bottom of the report.

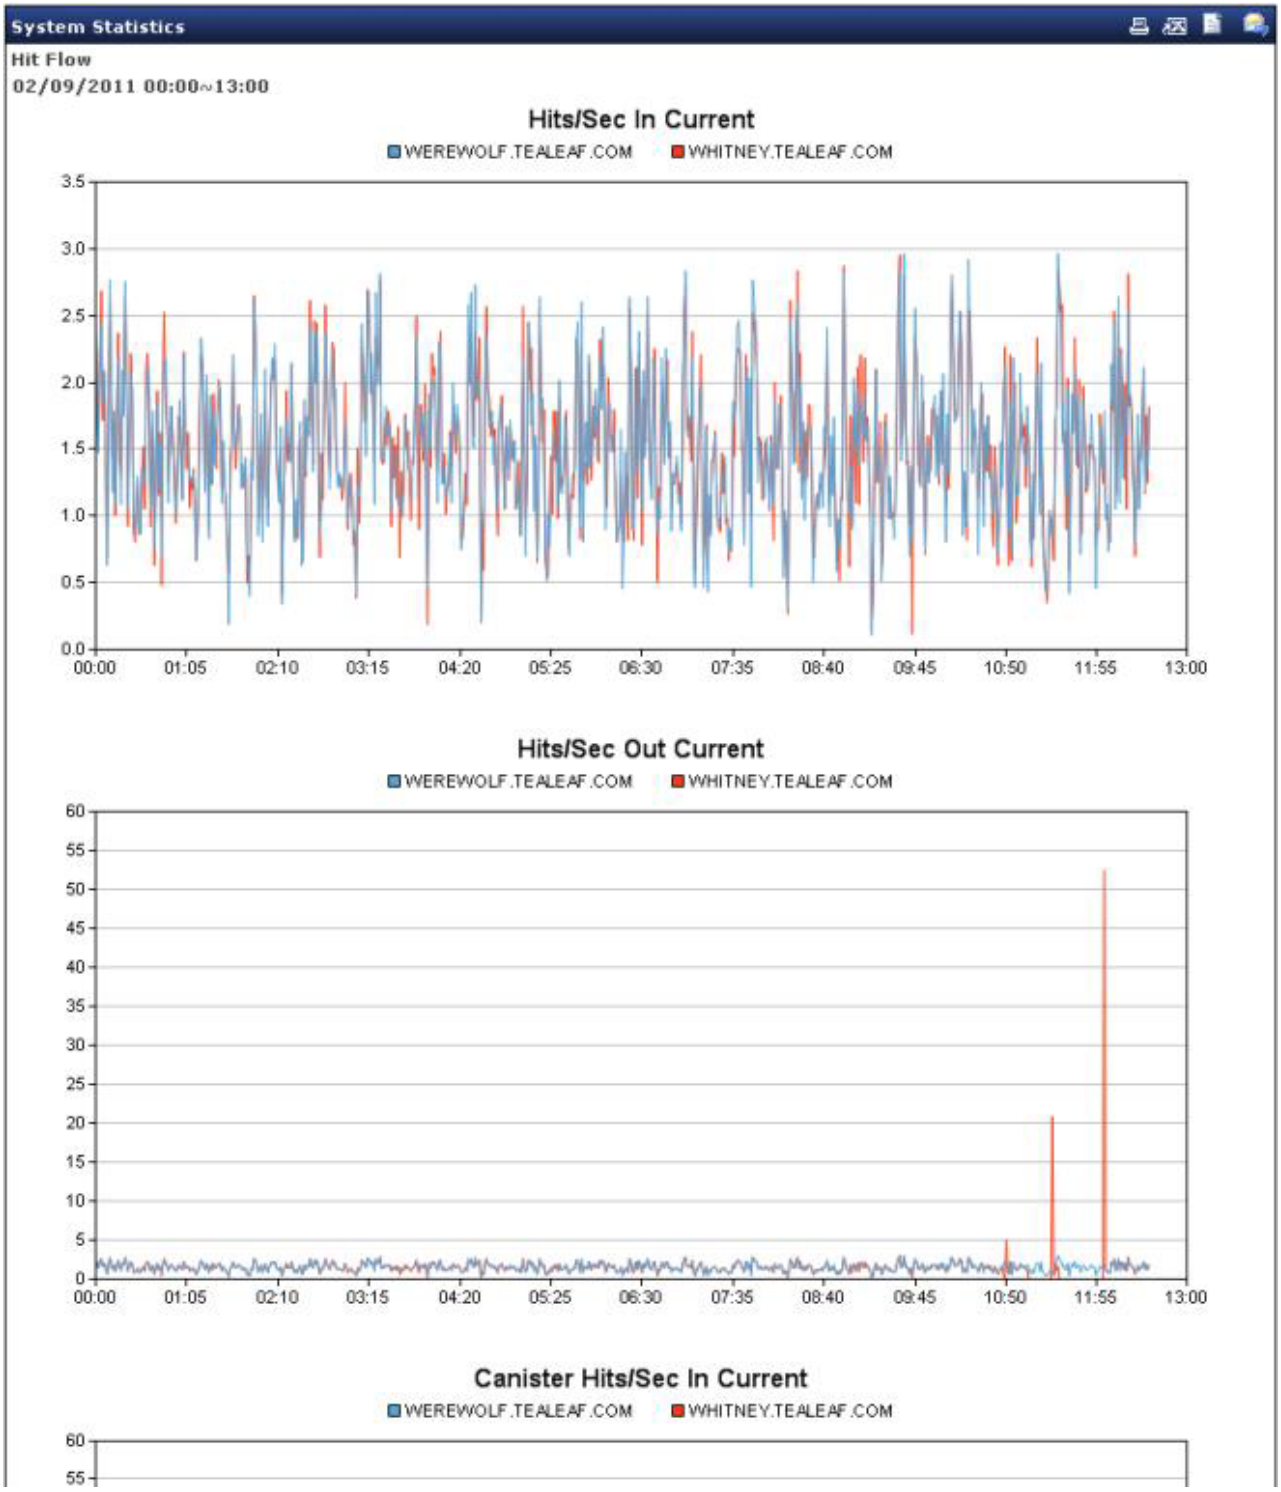

Figure 24. Custom Report - Hit Flow

Byte Flow

Provides volume in bytes of traffic flow in and out of each canister in the system over recent periods.

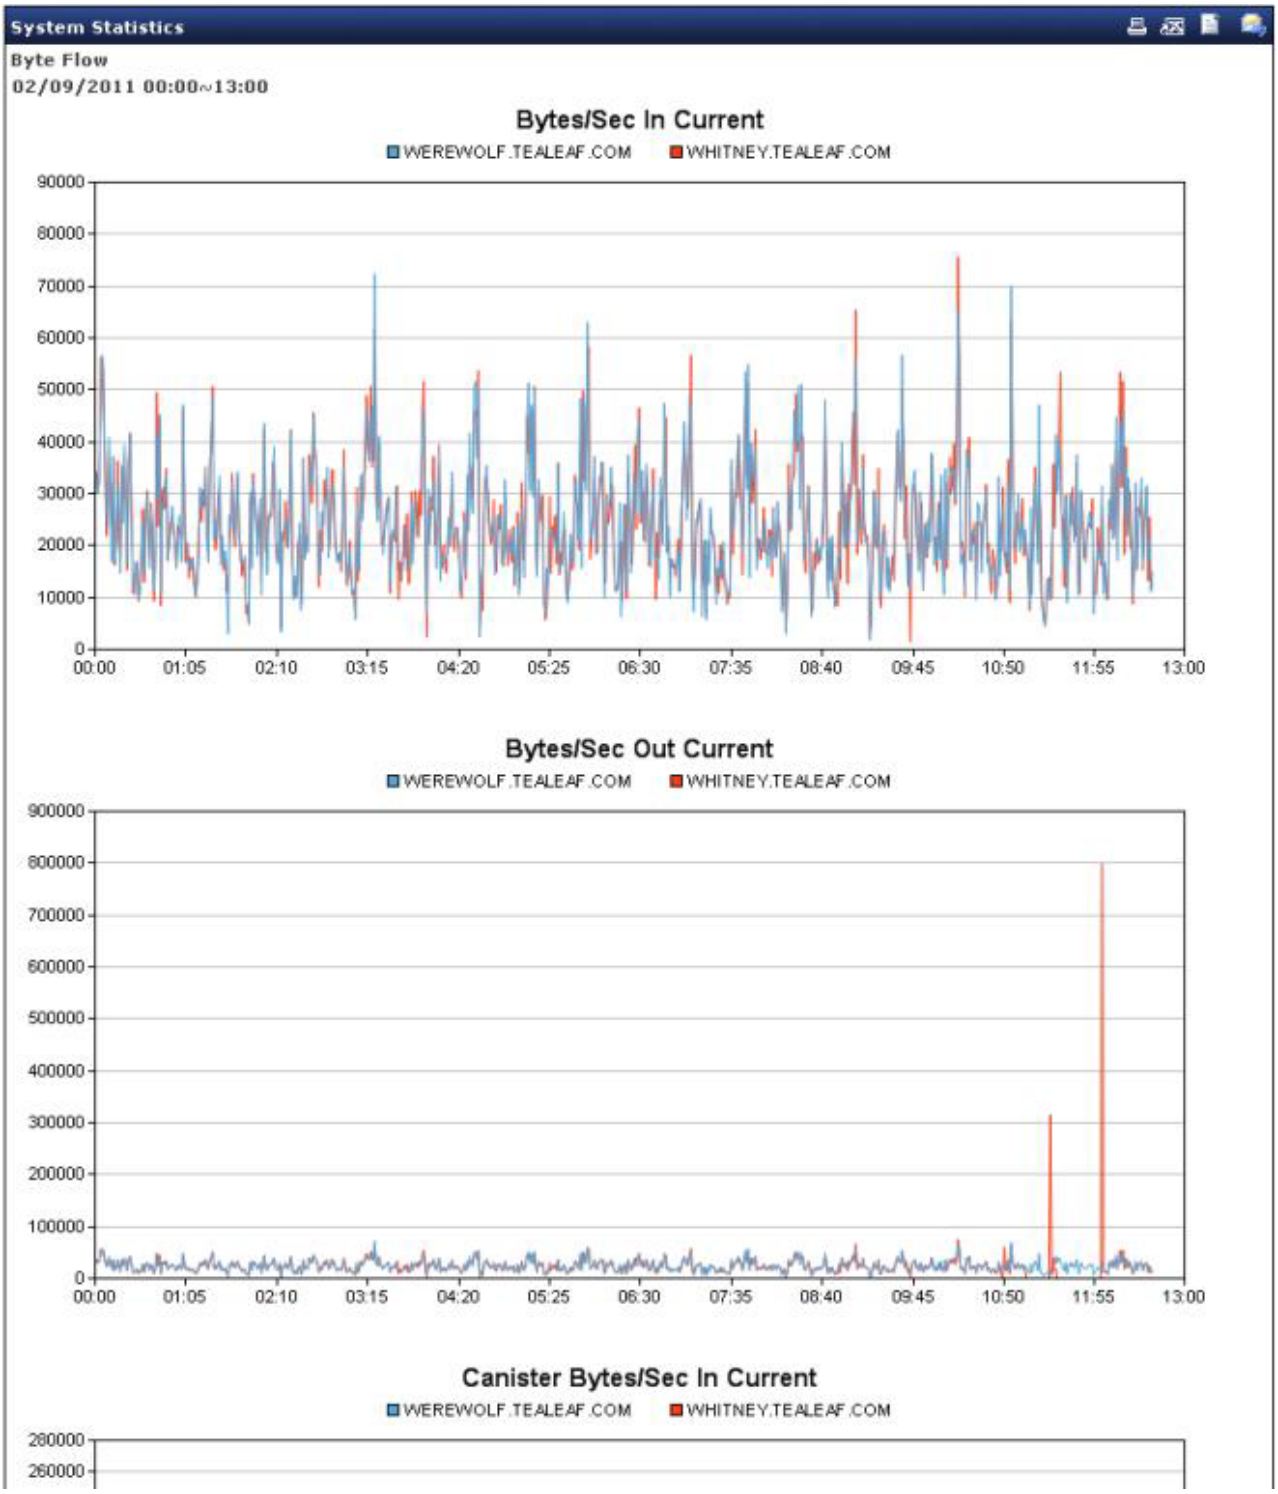

Figure 25. Custom Report - Byte Flow

Queuing

Summarizes queuing information for each canister.

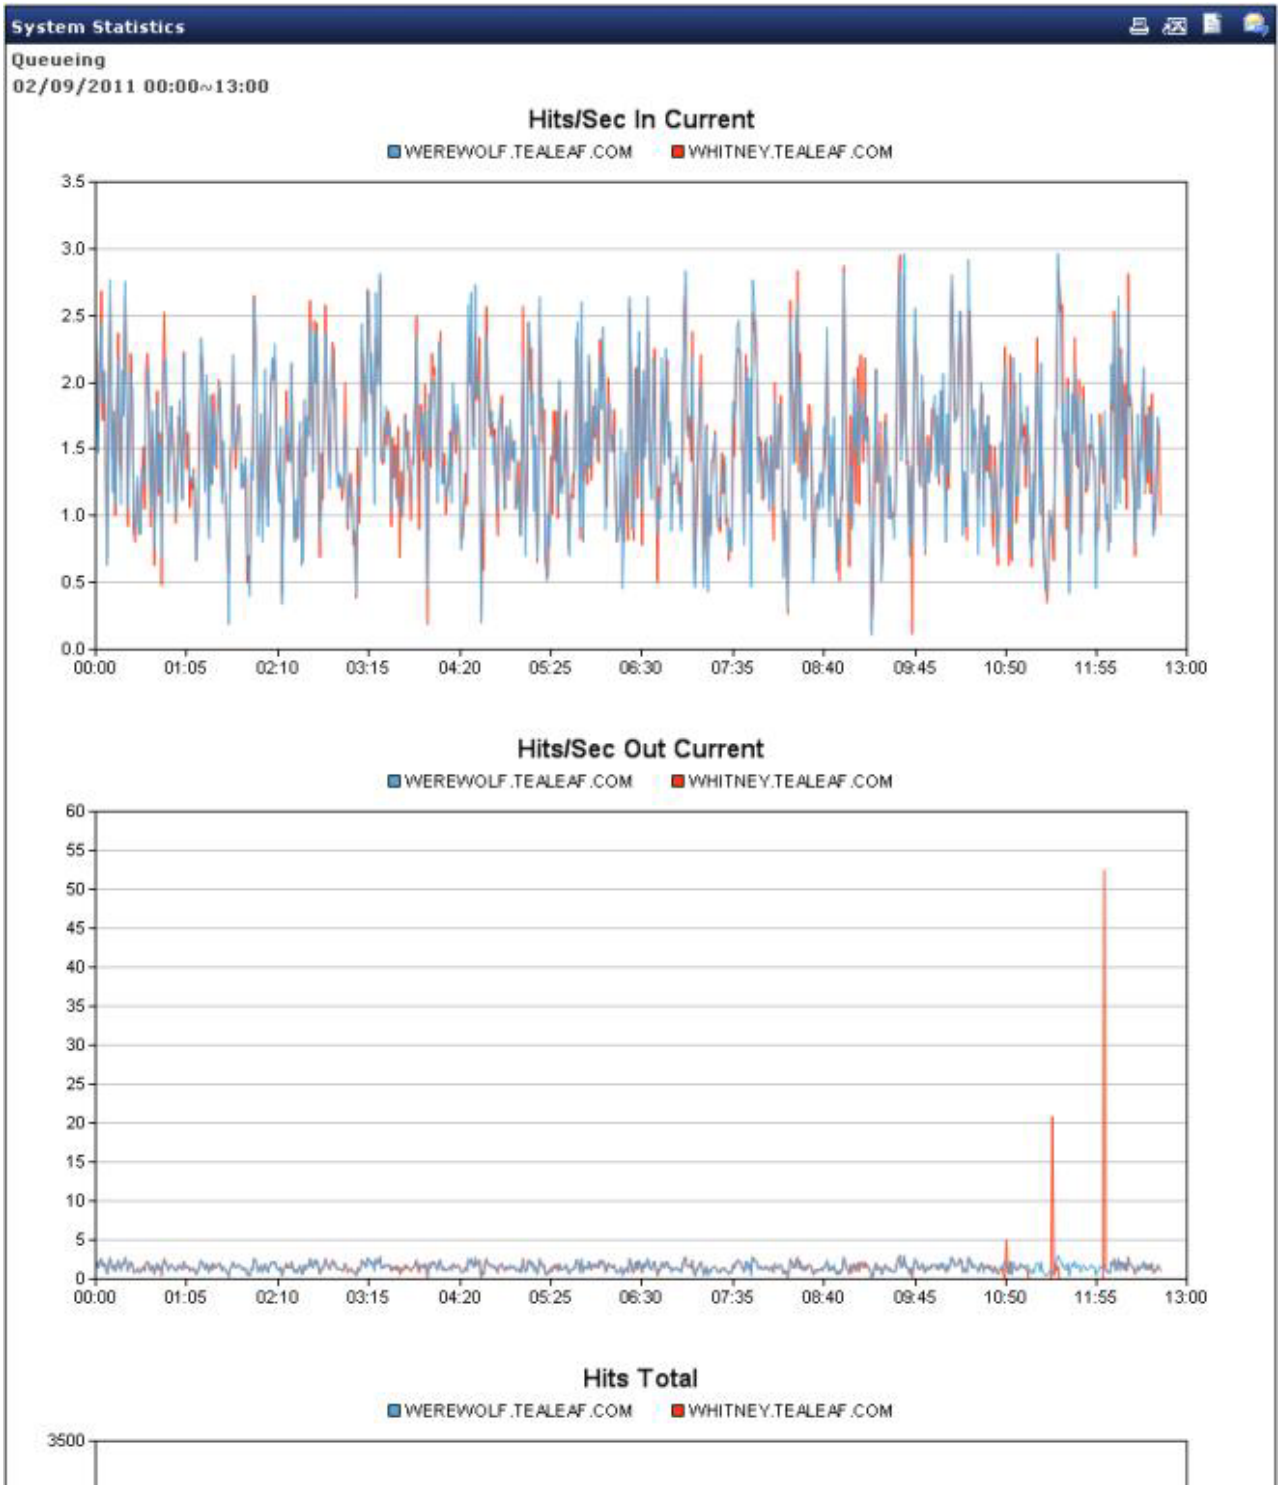

Figure 26. Custom Report - Queueing

CX Capture Hits

Identifies the volume of hits of non-SSL traffic, flows to delivery destinations, and current volume per second.

- Delivery destinations can be defined through the PCA Web Console. See "Passive Capture Configuration via Web Console" in the *IBM Tealeaf Passive Capture Application Manual*.

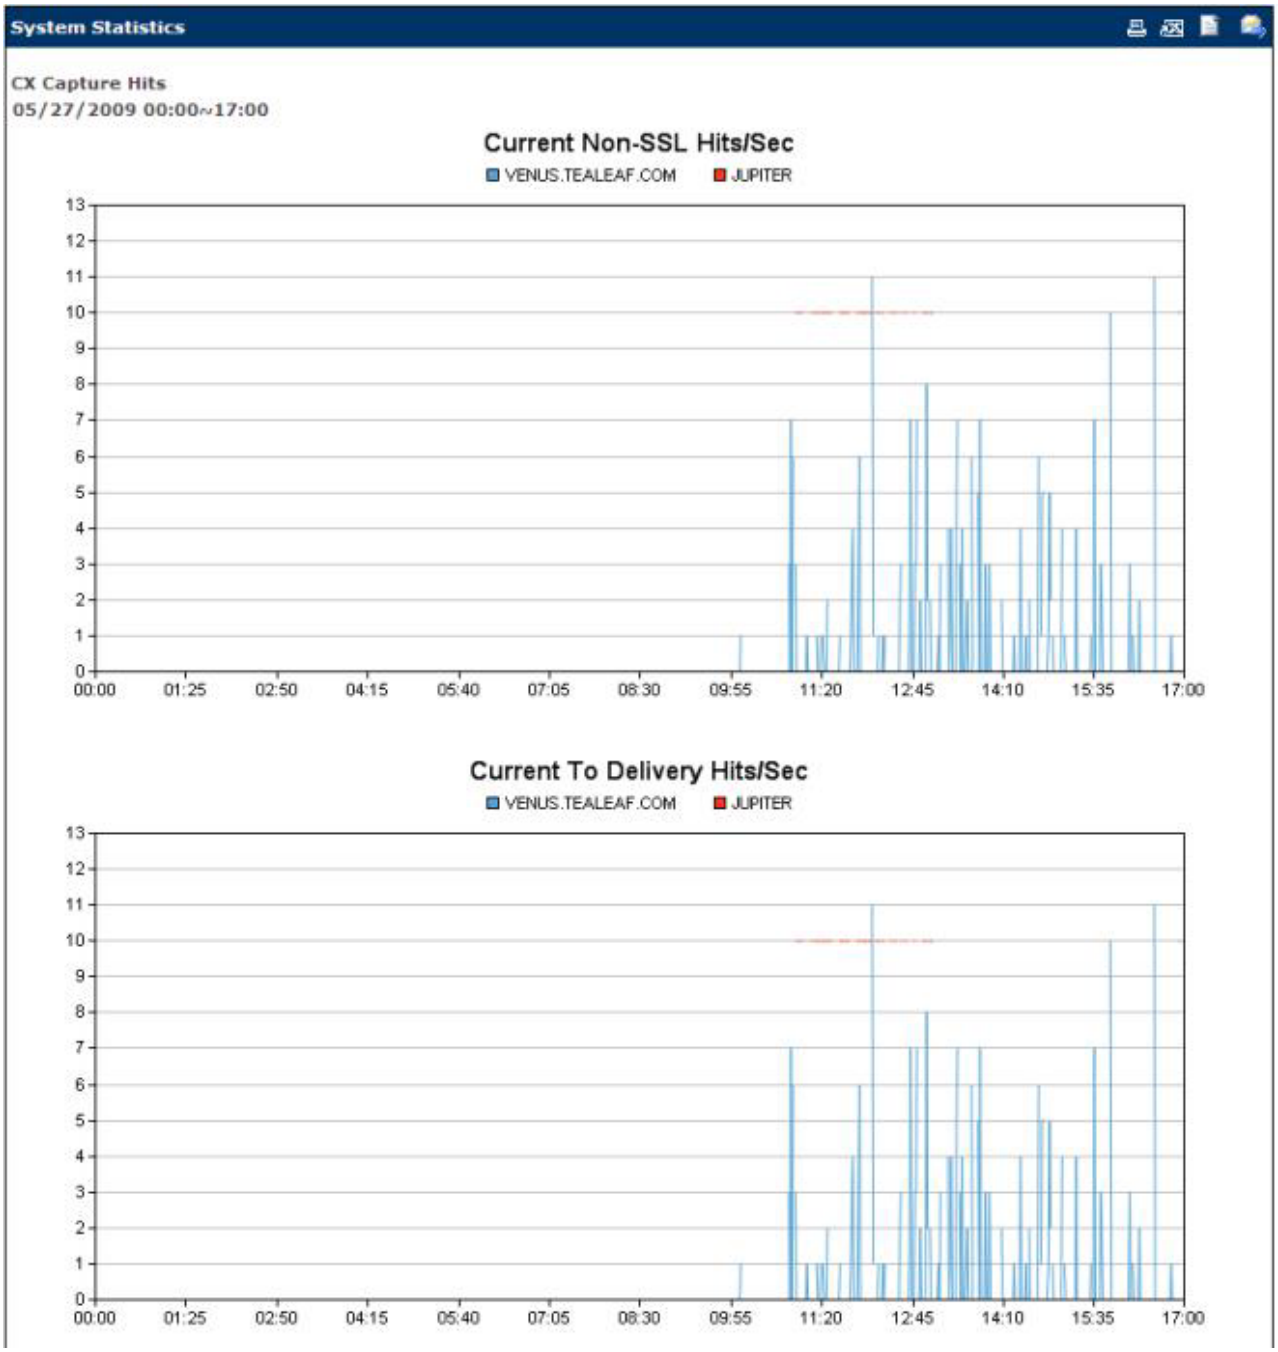

Figure 27. Custom Report - IBM Tealeaf CX Capture Hits

CX Capture Memory

Indicates memory usage by each process in the IBM Tealeaf CX Passive Capture Application over the reporting period.

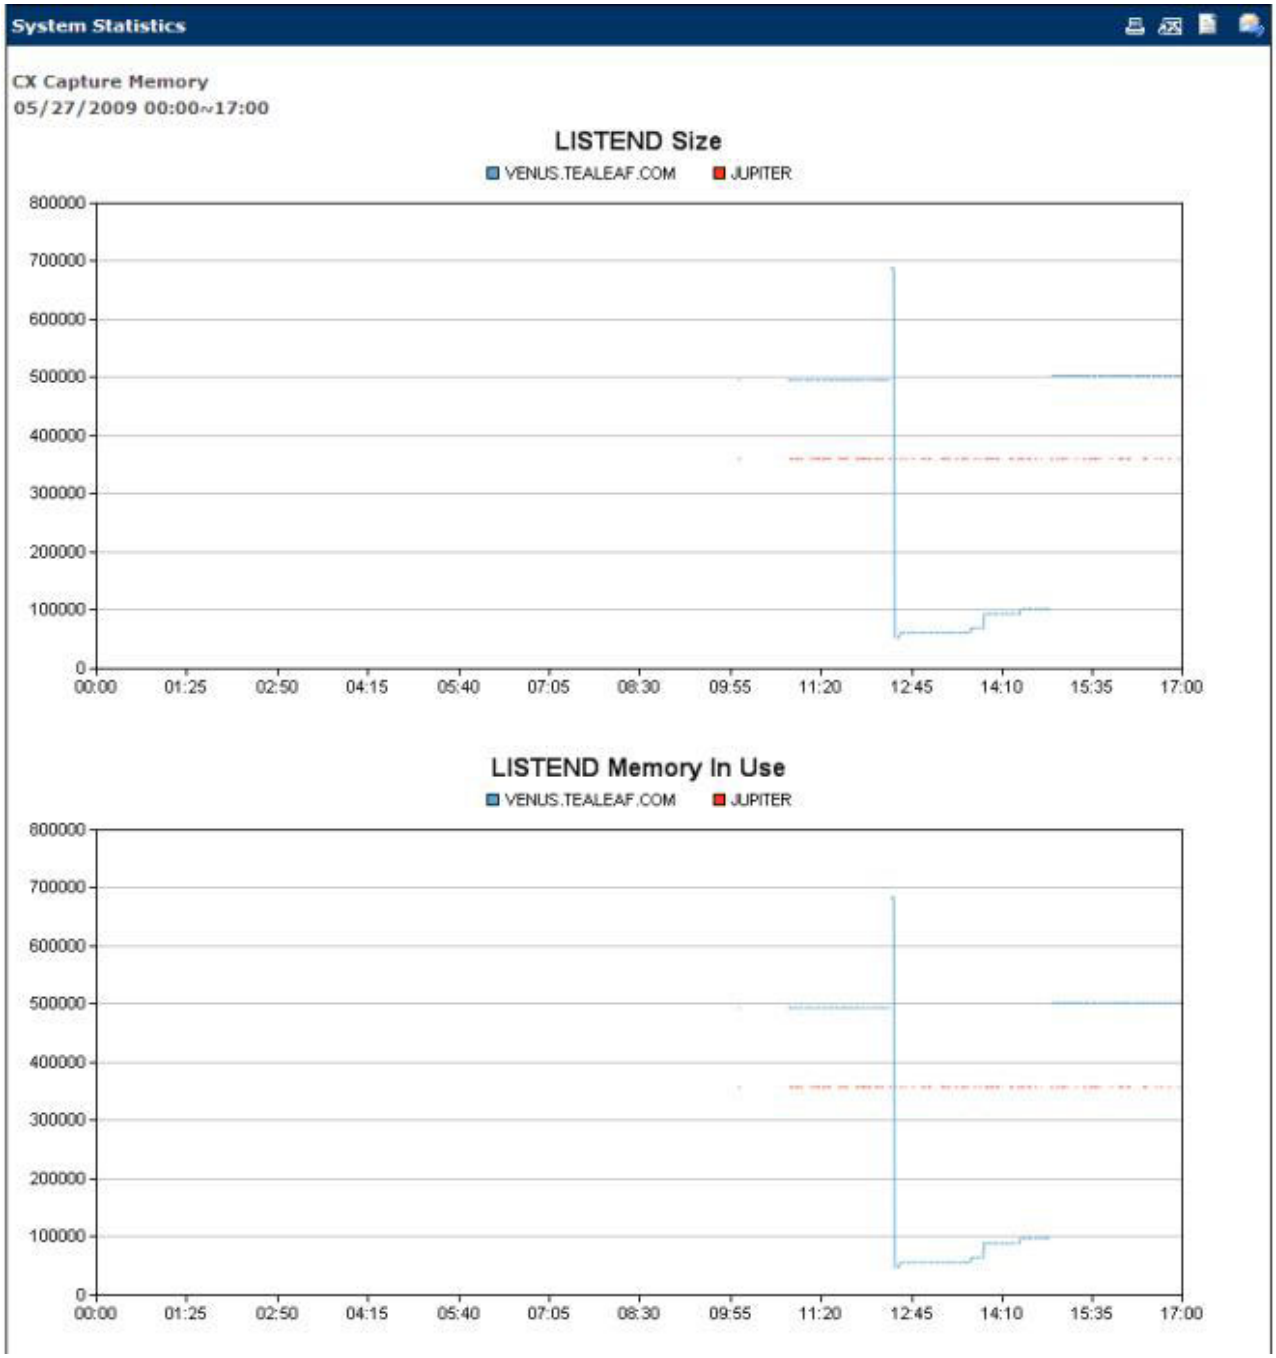

Figure 28. Custom Report - IBM Tealeaf CX Capture Memory

CX Capture Queue

This report indicates the volume of hits that have been queued for delivery by the PCA over the time period.

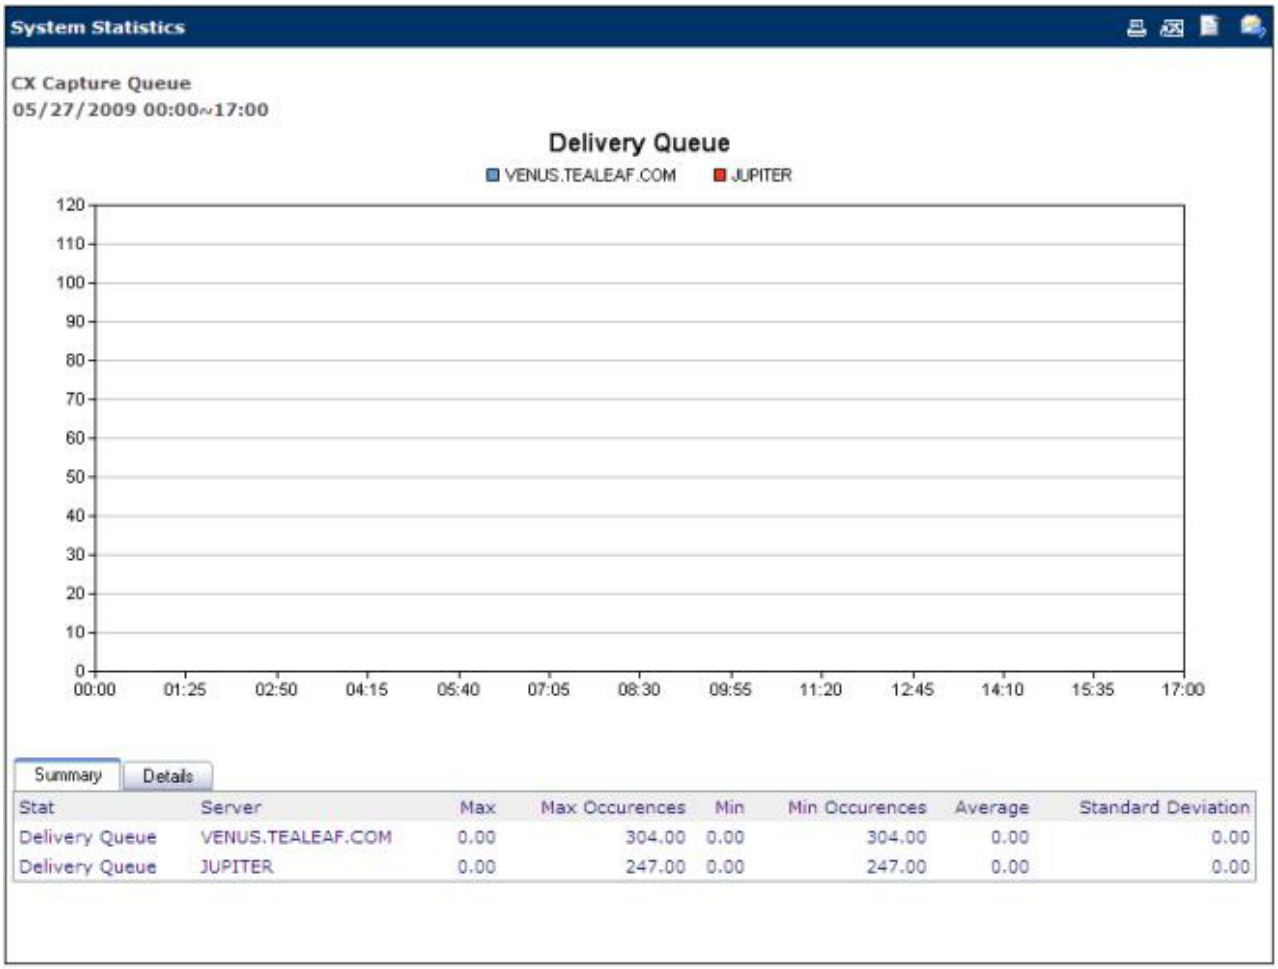

Figure 29. Custom Report - IBM Tealeaf CX Capture Queue

CX CPU

Indicates the CPU load for each capture server in the system.

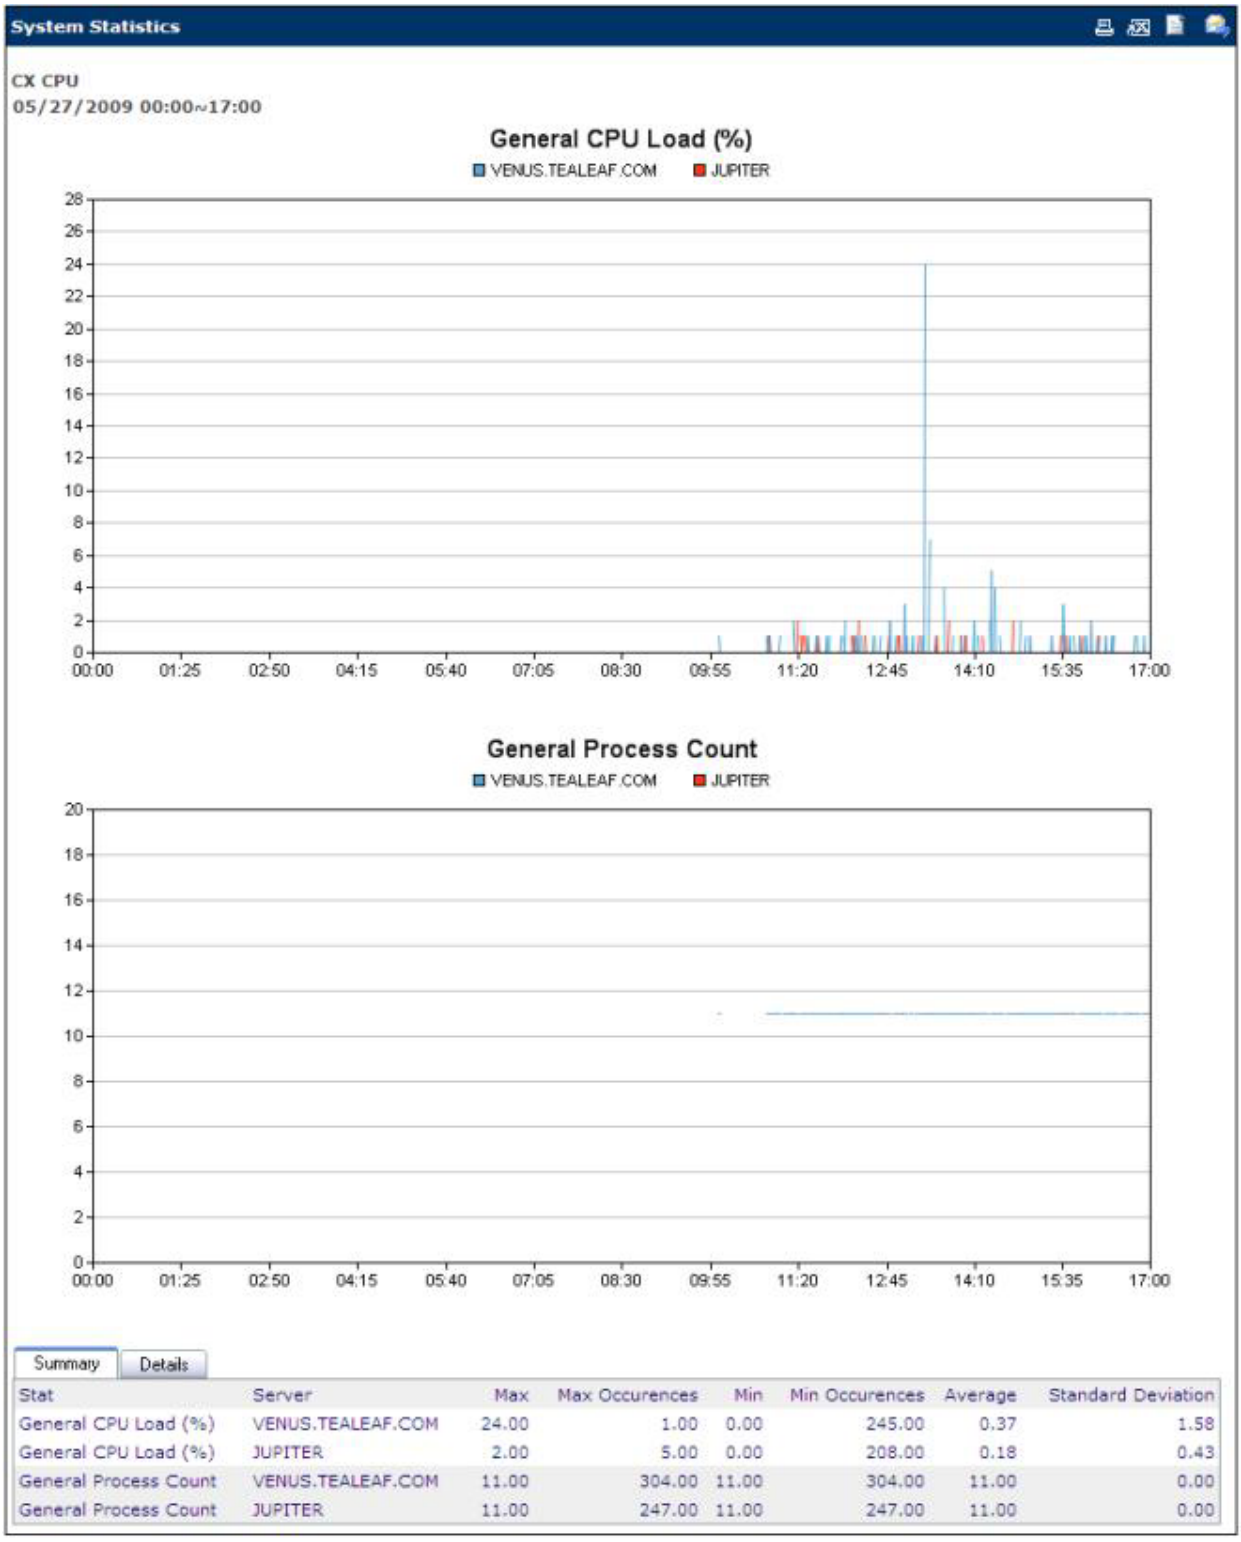

Stat	Server	Max	Max Occurences	Min	Min Occurences	Average	Standard Deviation
General CPU Load (%)	VENUS.TEALEAF.COM	24.00	1.00	0.00	245.00	0.37	1.58
General CPU Load (%)	JUPITER	2.00	5.00	0.00	208.00	0.18	0.43
General Process Count	VENUS.TEALEAF.COM	11.00	304.00	11.00	304.00	11.00	0.00
General Process Count	JUPITER	11.00	247.00	11.00	247.00	11.00	0.00

Figure 30. Custom Report - IBM Tealeaf CX CPU

Active Sessions

Identifies the number of active and total sessions in the Short-Term Canister over the selected time period.

- Minimum, maximum, and average computations are available at the bottom of the report.

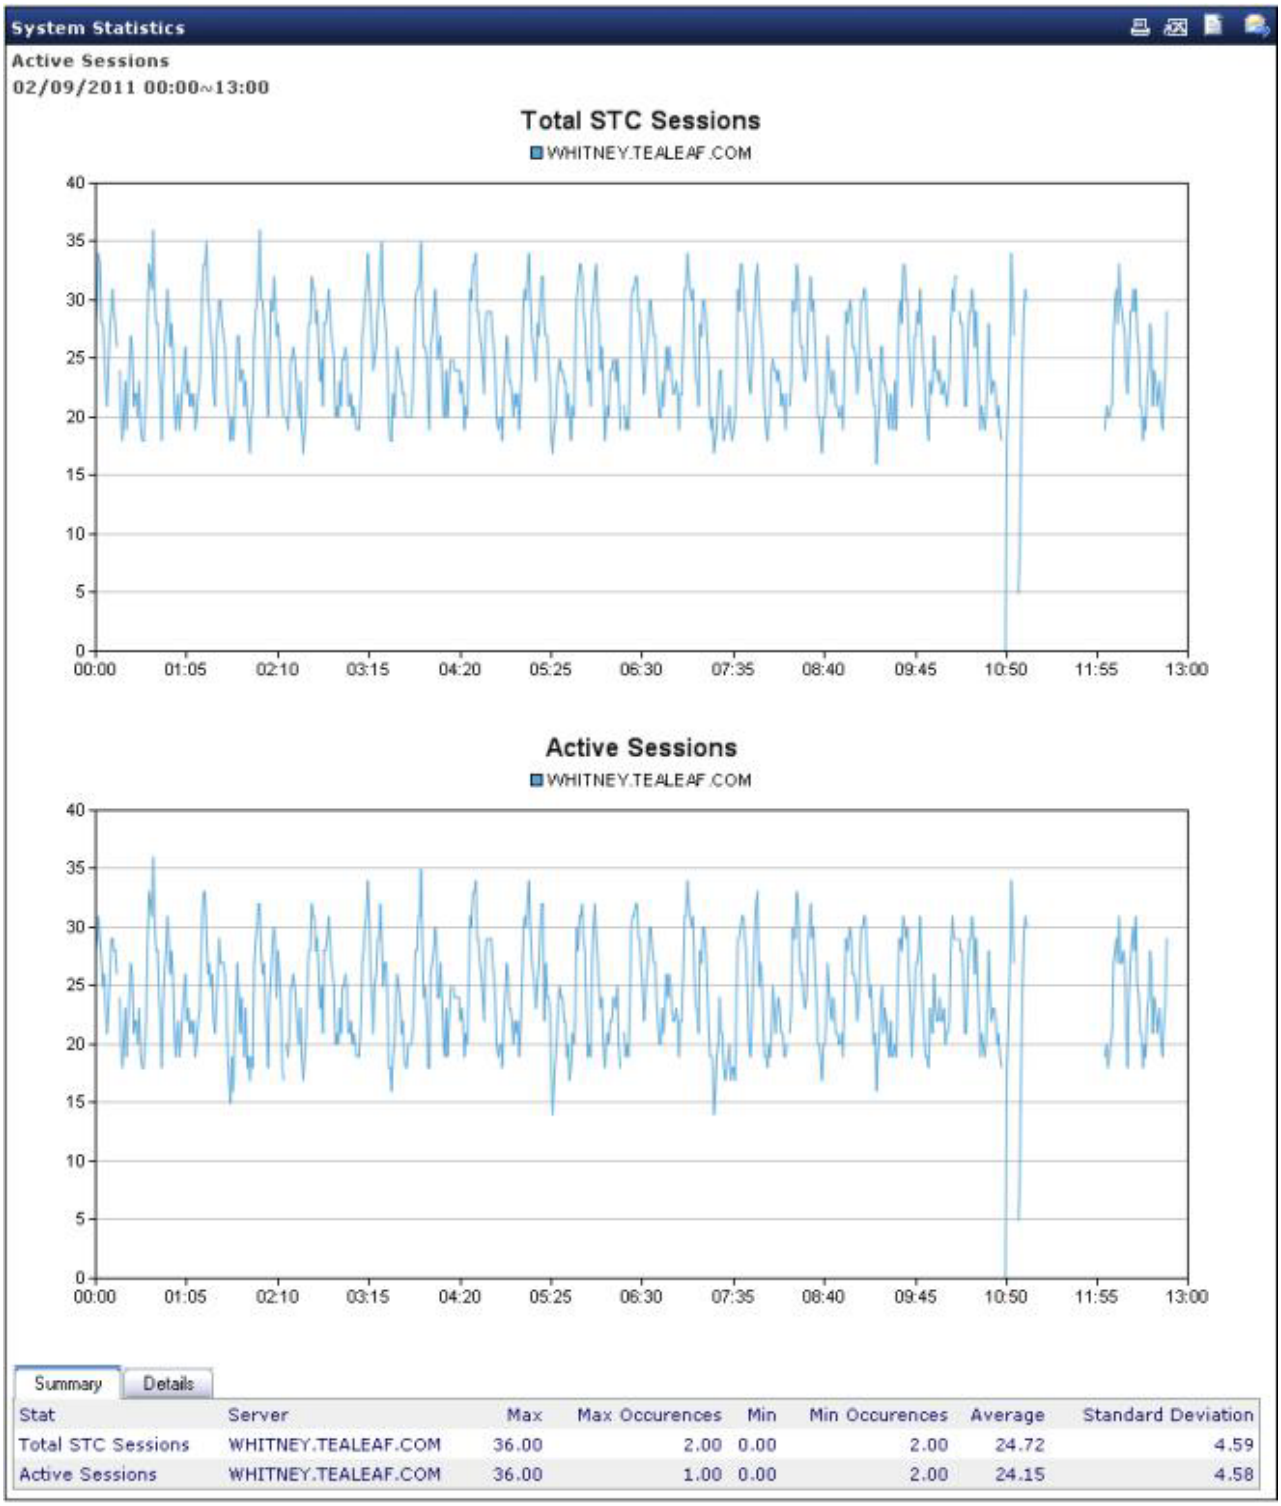

Figure 31. Custom Report - Active Sessions

Canister Reporting Data

This report indicates the volume of records written from the canister to the appropriate database(s) for events and dimensions.

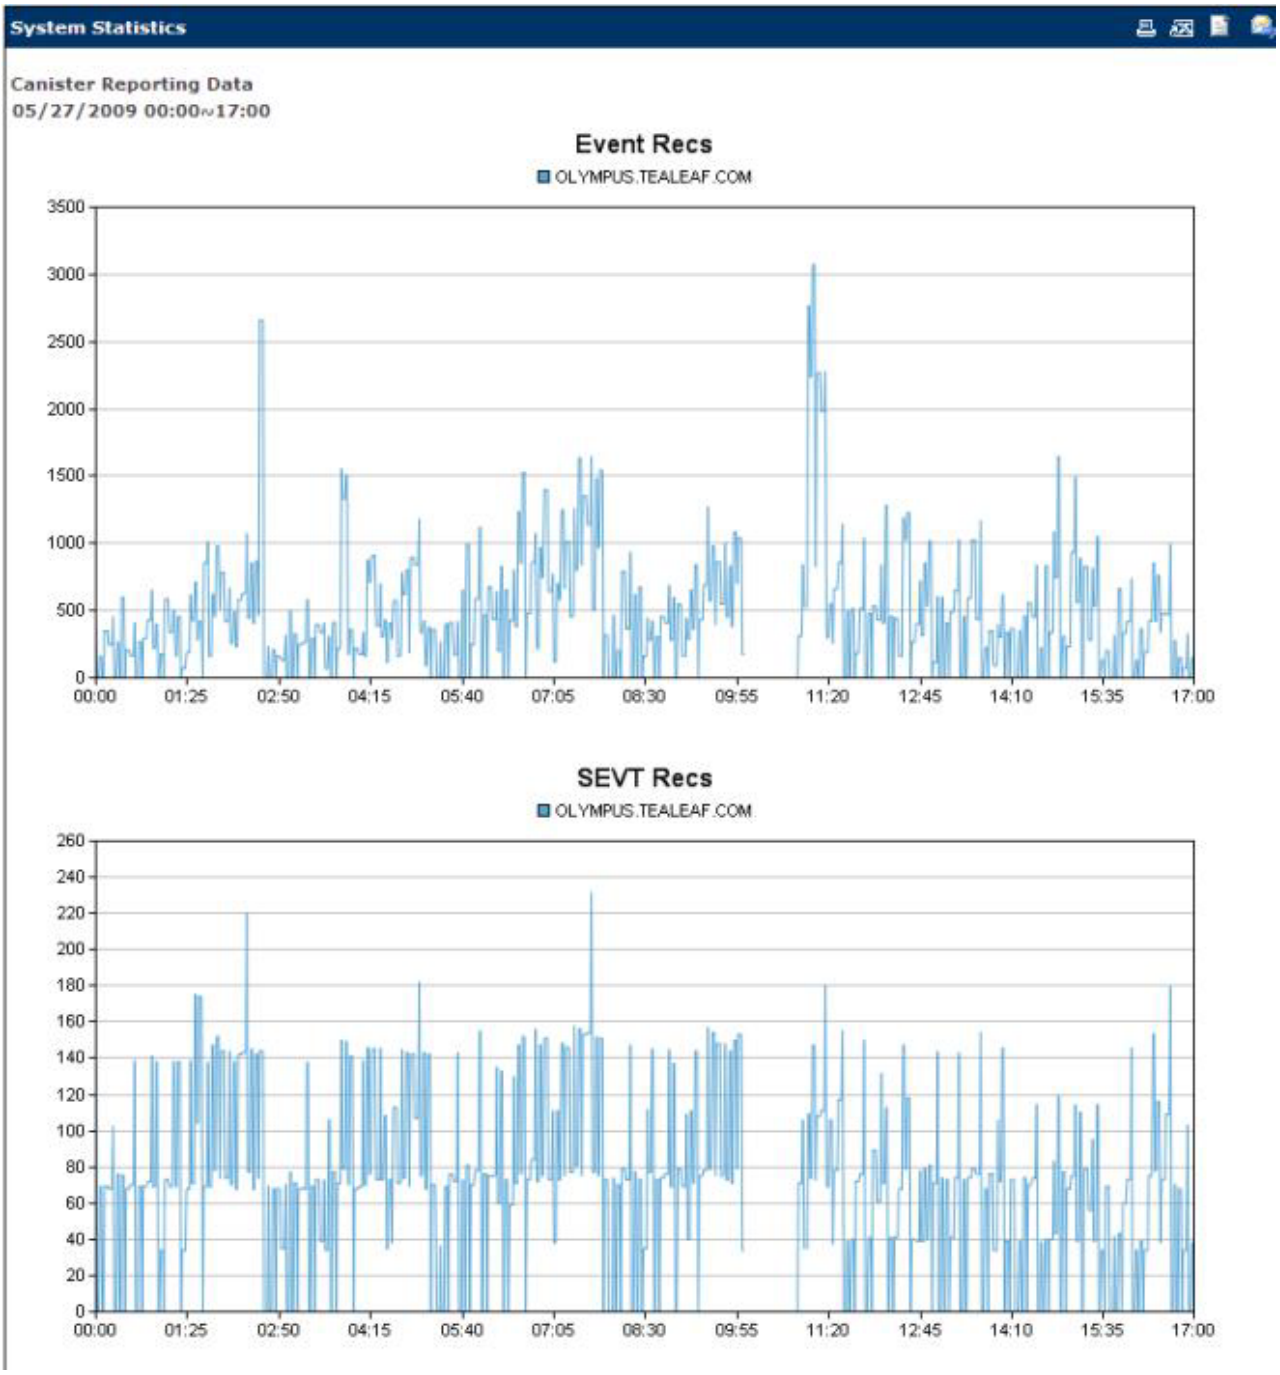

Figure 32. Custom Report - Canister Reporting Data

Indexing

Over the reporting period, this report indicates the volume of sessions that could not be indexed.

Figure 33. Custom Report - Indexing

STC Performance

Provides compression information of the Short-Term Canister over the reporting period.

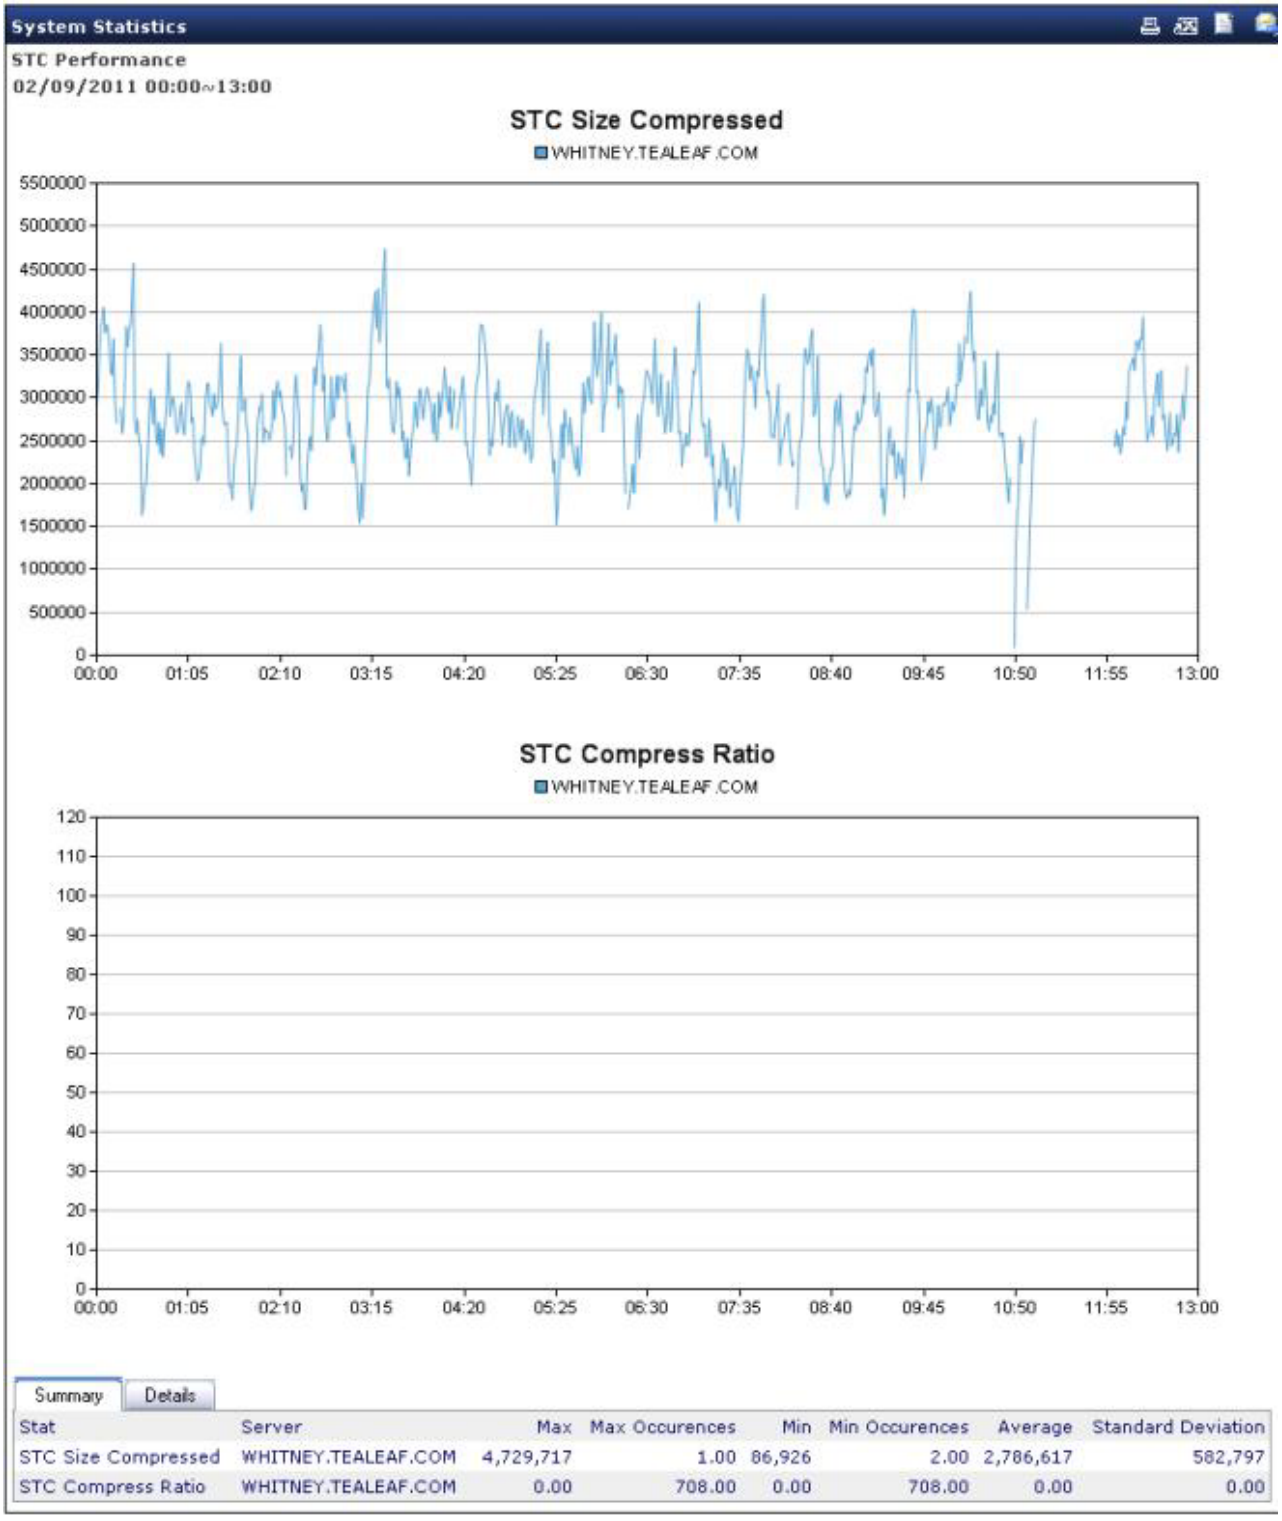

Stat	Server	Max	Max Occurences	Min	Min Occurences	Average	Standard Deviation
STC Size Compressed	WHITNEY.TEALEAF.COM	4,729,717	1.00	86,926	2.00	2,786,617	582,797
STC Compress Ratio	WHITNEY.TEALEAF.COM	0.00	708.00	0.00	708.00	0.00	0.00

Figure 34. Custom Report - STC Performance

Canister Memory

Indicates the memory usage over time for the CTree canister.

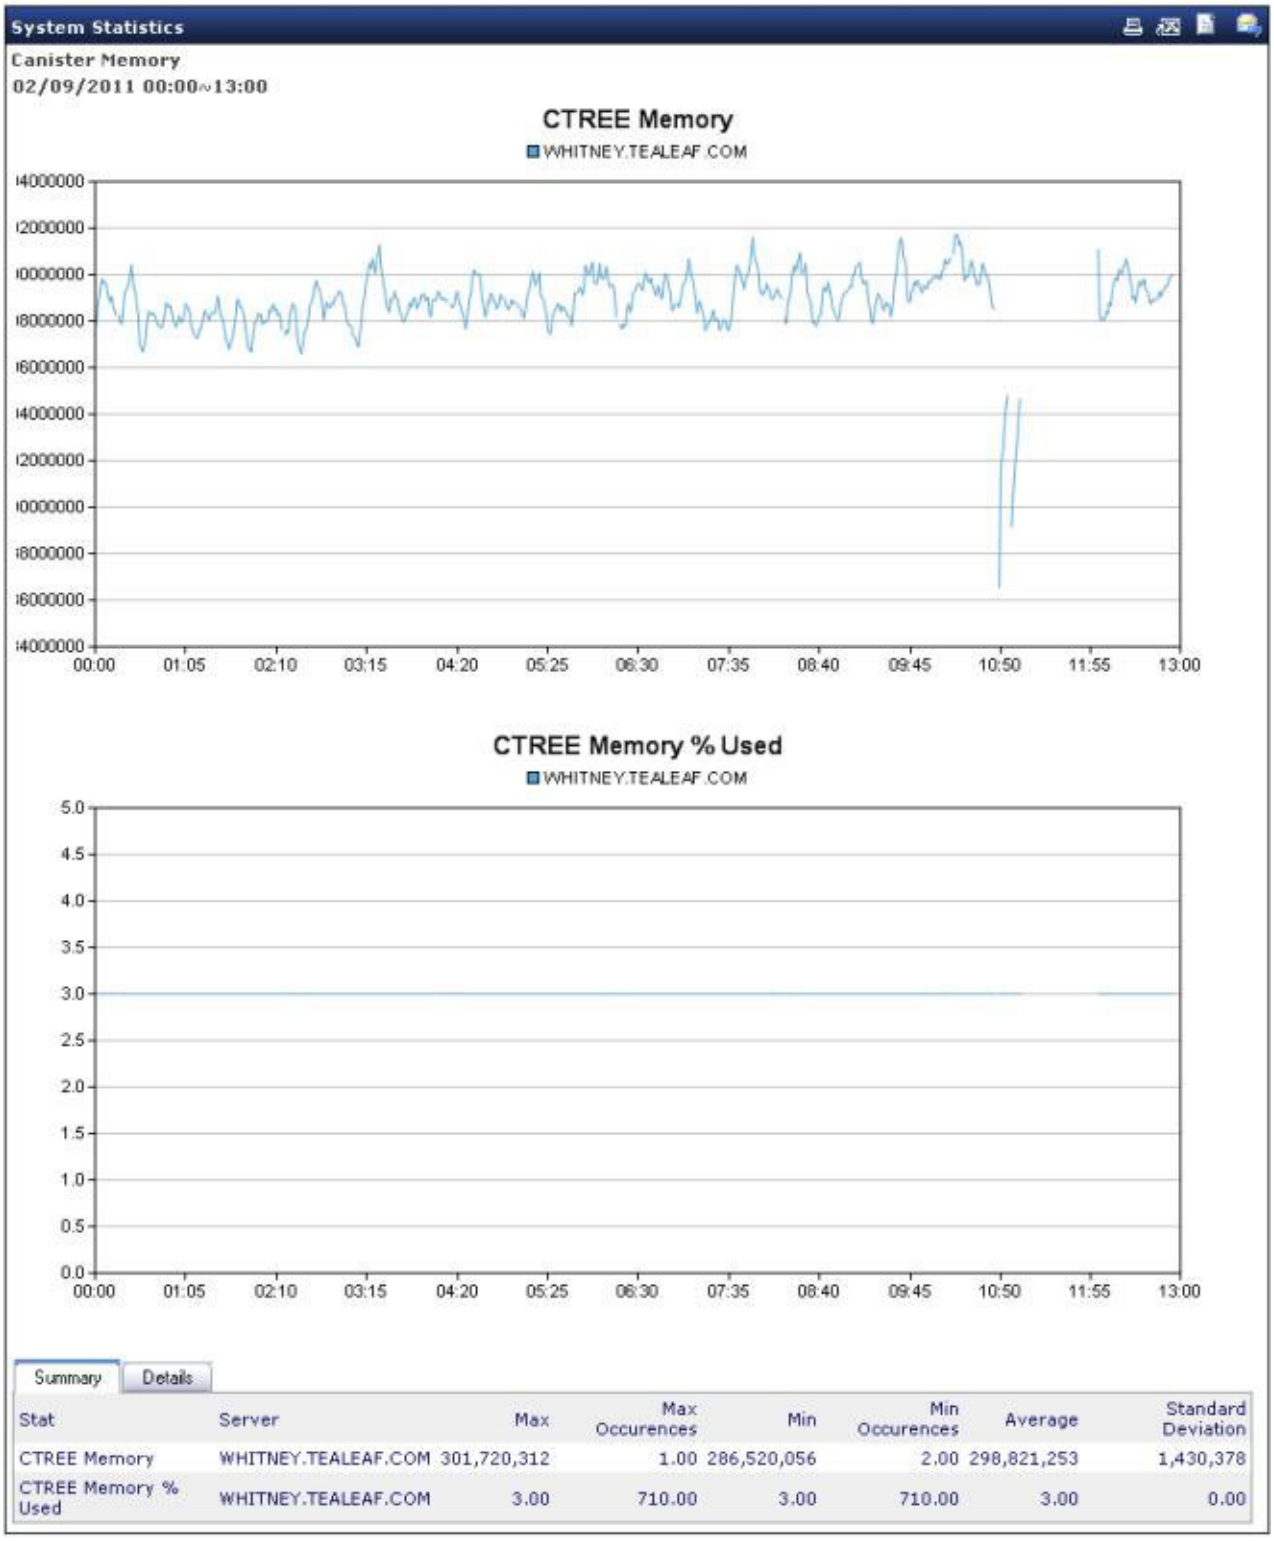

Stat	Server	Max	Max Occurences	Min	Min Occurences	Average	Standard Deviation
CTREE Memory	WHITNEY.TEALEAF.COM	301,720,312	1.00	286,520,056	2.00	298,821,253	1,430,378
CTREE Memory % Used	WHITNEY.TEALEAF.COM	3.00	710.00	3.00	710.00	3.00	0.00

Figure 35. Custom Report - Canister Memory

Canister CPU

Canister CPU usage over recent time periods.

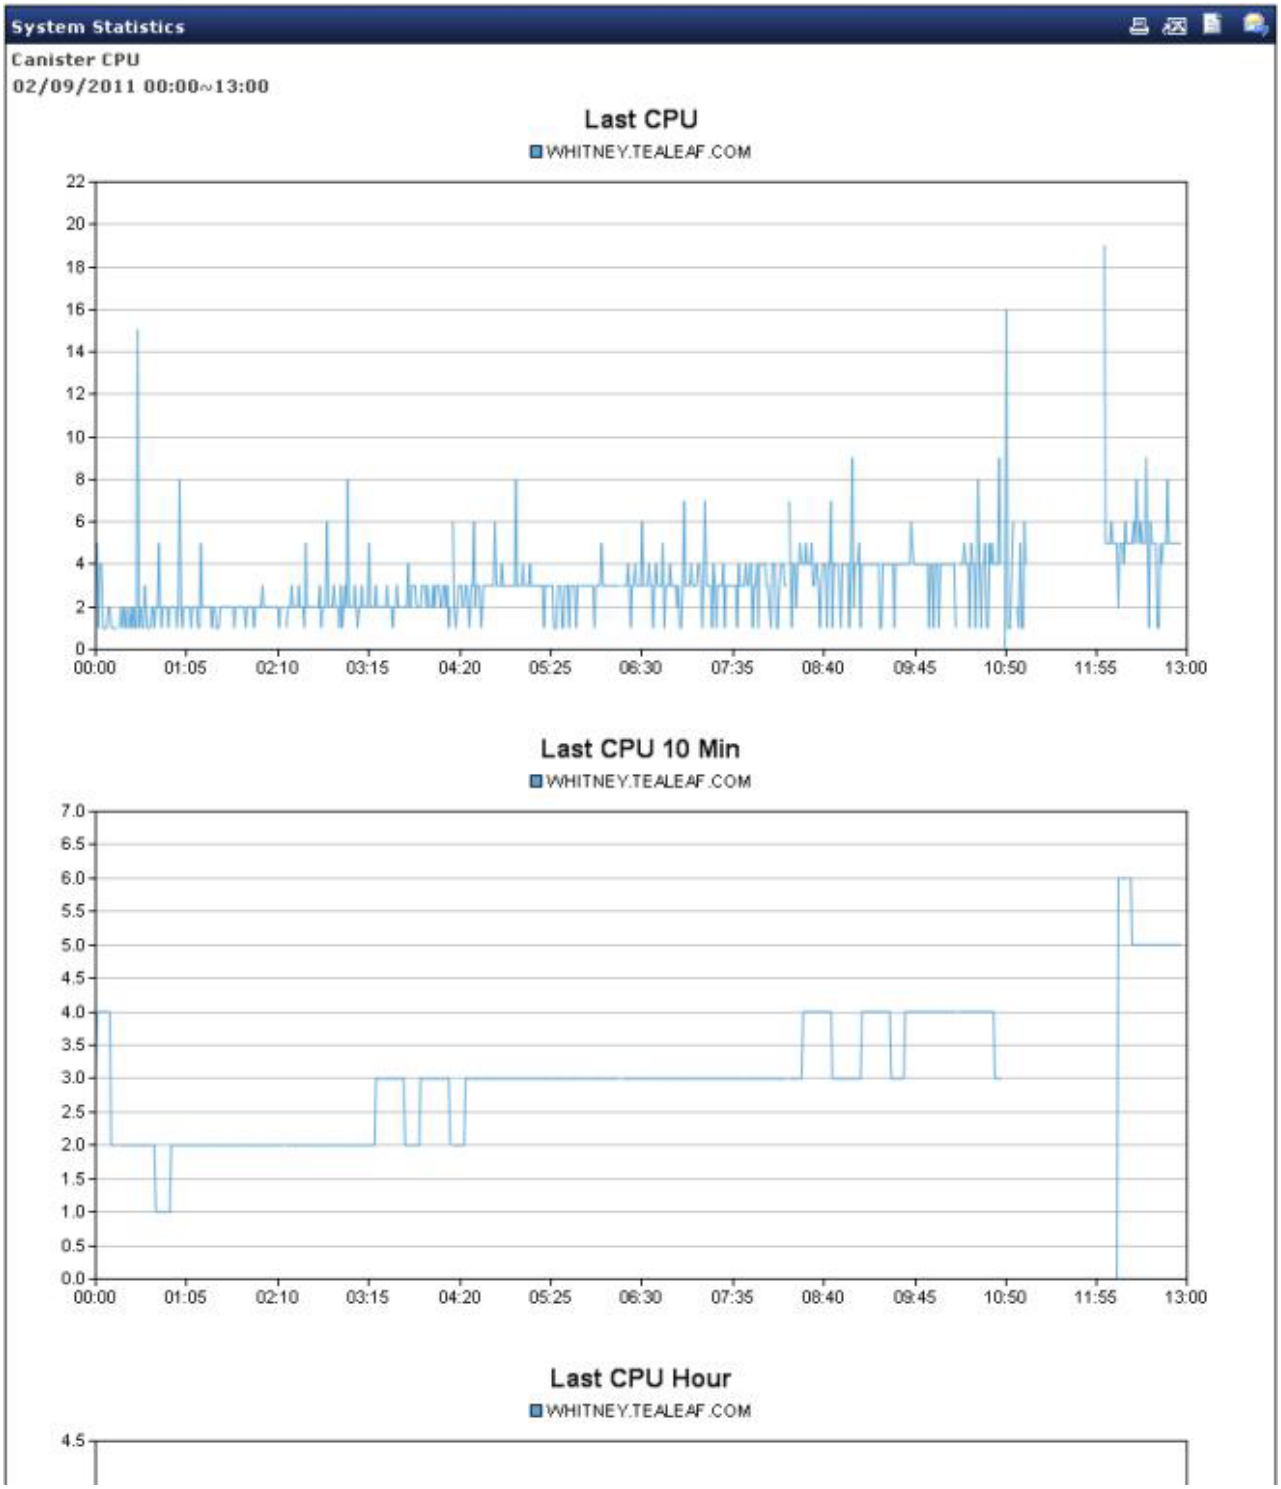

Figure 36. Custom Report - Canister CPU

Canister Statistics

By default, the Tealeaf Canister monitors statistics and periodically submits statistics hits for capture and storage in the database. These statistics are available for display in this report in the **Canister stats** category.

- These statistics are also available in the System Status page. See "System Status" on page 8.

Summary Tab: In the Summary tab, you can review statistical calculations, such as minimum, maximum, average, and standard deviation for each of the above statistics.

Details Tab: In the Details tab, you can breakdowns of each statistic for each hour in the report.

Monitoring User Activity

Through the Portal, you can keep tabs on user activities through a set of filtered reports. These reports provide insight into Tealeaf users who have logged in or have not logged in for a period of time. Additional usage and session review reports are available.

Note: The Portal application automatically logs all activities by user ID, including session search and session access for replay purposes. For RTV, only replays are audited per user. This auditing cannot be disabled.

Configuring User Activity Reports

1. From the Portal menu, select **Tealeaf** > **User Activity**.
2. Dashboard versions of the report are displayed in each tab. Click the appropriate tab.
   - "Logs tab" on page 63
   - "Activity Counts tab" on page 65
   - "Activity by User/Group tab" on page 67
3. To open an individual report, click the Open icon in the toolbar.
   - See "Toolbar" on page 58.
4. Configure the date range. See "Date Range Selector" on page 58.
5. User Activity reports can be filtered based on user, group, or dimension value. See "Options" on page 60.
6. To refresh the data, click the Refresh icon in the toolbar.
7. Depending on the report and your user settings, you may click links in the report to reveal detail reports on the displayed data.
8. To export the report data for local use, click the Export icon in the toolbar. See "Export" on page 77.

Toolbar

The following toolbar is displayed above each report.

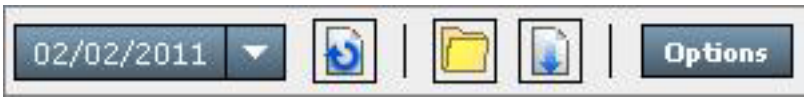

Figure 37. User Activity reports toolbar

Table 2. Toolbar

Icon	Command	Description
	Date Selector	Click to select a new date range to display in the report. See "Date Range Selector."
	Refresh	Refresh the report data.
	Open	Open and display a different User Activity report.
	Export	Export report data. See "Export" on page 77.
	Options	Change User Activity report options. See "Options" on page 60.

Date Range Selector

You may select a single day or range of dates to display in the report.

- To open the Date Range Selector, click the date displayed in the toolbar.

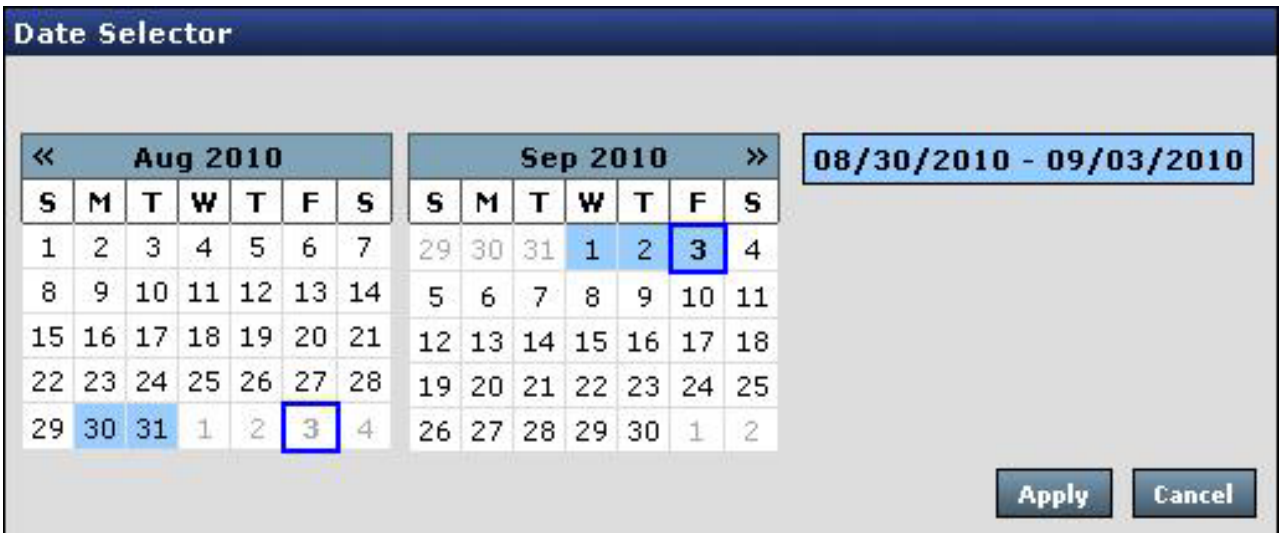

Figure 38. Performance Reports - Date Range Selector

In the above, the business week of 8/30/10 to 9/3/10 is selected for display.

- To configure a different date, click the date, or you may use the arrow buttons on each month to navigate to different months.
- To select a range of dates, click the earlier date. Then, click the latter end date. The range of dates in between and inclusive is selected.
- To apply the new date or date range to the report, click **Apply**. The report is automatically updated.
- To cancel changes to the date range, click **Cancel**.

Options

You may configure the following options for the displayed User Activity report.

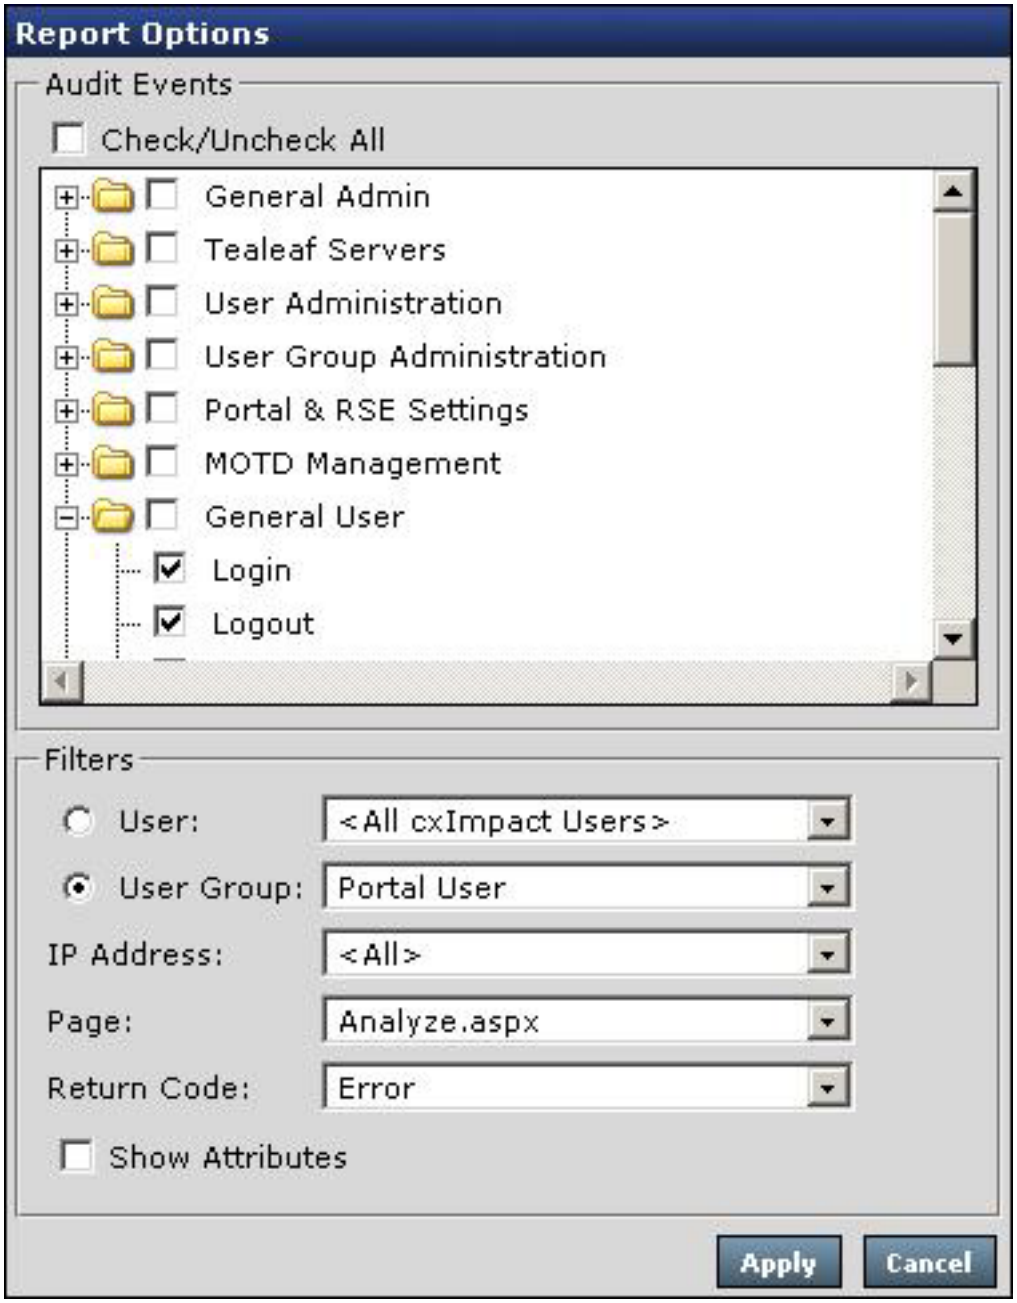

Figure 39. User Activity Reports - Options

The following options are available.

- When one or more of these filters is applied, they are listed at the top of the updated report.

Option Description

Audit events

Select the events to appear in the report. See "Audit events" on page 61.

Filters Apply user and dimension filters to your report.

- See “Filtering by user or group.”
- See “Applying a dimension filter.”

- To apply your changes to the report options, click **Apply**.
- To cancel changes, click **Cancel**.

Audit events

You can specify the user events for which you wish to see an audit report.

- Expand audit event groups and select the events you wish to review.
- To toggle selection of all audit events, click the Check/Uncheck All checkbox.

In the report panel, each selected audit event appears as a different data series in the report.

Filtering by user or group

The audit reports let you analyze site visitor and Tealeaf Portal activity, filtering the output by user or group members or by dimension. Events can be filtered by the following.

- These options persist if you load a different report.
- For more information on filtering by dimension, see “Applying a dimension filter.”

Filter **Description**

User Filter the report by all users of a Tealeaf product or by individual users.

User Group
Filter the report by specific user groups or by all user groups of a specified Tealeaf product.

Applying a dimension filter

Optionally, you can filter any User Activity report to display data for a selected value from one or more of the following dimensions. You may filter the report for one specific value per dimension. You may use any number of the listed dimensions to filter your report.

Dimension
Description

`IP Address`
Filter entries by visitor IP address.

`Page` Filter entries to a specific Portal page.

`Return Code`
Filter entries to those with Success or Error codes.

- `0` - success
- `1` - error

- To display all of the attributes and details associated with the event, click the Show Attributes checkbox.

The IP Address and Page values are captured in a default report group provided by Tealeaf as part of the installation. See "TEM Dimensions Tab" in the *IBM Tealeaf Event Manager Manual*.

Common report fields

The fields in the following table appear in one or more of the user activity reports.

Field **Description**

User For user-based reports, this field indicates the Tealeaf user name to which the row data applies.

Group For group-based reports, this field indicates the Tealeaf name of the group to which the row data applies.

Searches
For the focus period, this column indicates the number of searches executed by the user or group.

Sessions Returned
Number of session returned from search.

Sessions Replayed
Number of returned sessions that were replayed.

Search Effectiveness %
Percentage of searches that returned non-zero results. Higher percentages indicate a better ability to specify search terms.

Search Efficiency %
This percentage is the inverse of the average number of session returned, when that number is not zero. Higher percentages of search efficiency indicate that the user is better at specifying searches that return a limited number of results, enhancing efficiency in locating data.

Replay Use %
The number of replayed sessions divided by the number of returned sessions. This figure can be greater than 100% since a single returned session can be replayed multiple times.

Last Login
Timestamp for the last login by the user or group.

- Times are based on the Tealeaf system time zone.

Last Search
Timestamp for the last search by the user or group.

- Times are based on the Tealeaf system time zone.

Last Replay
Timestamp for the last replay by the user or group.

- Times are based on the Tealeaf system time zone.

Audit Reports

The following reports are available.

- Metrics which tabulate the number of searches per user or user group include counts of searches initiated from the Tealeaf Portal and the RTV desktop application.
- For more information on RTV, see "RealiTea Viewer (RTV) User Manual" in the *IBM Tealeaf RealiTea Viewer User Manual*.

User Activity Report

This section describes the user activity reports.

Logs tab

In the Logs tab, you can review for the selected users the logging information for the selected audit events for the focus period.

- To display only the logging information for a specific user, click one of the username links in the User column.
- To review session attribute information about the event, click the I icon in the left column of the report.
- By default, the number of days retained for the log is controlled by the Security Log Data- Days Retain system setting. To change this value, select **Tealeaf** > **Portal Management** > **IBM Tealeaf CX Settings** > **Data Retention**. See "CX Settings" on page 192.

Time Stamp	User	Page	Event	Return Code	Report Value
05/05/2010 13:58:10	ADMIN	Logout.aspx	Logout	Success	N/A
05/05/2010 13:57:16	ADMIN	Default.aspx	Login	Success	N/A
05/05/2010 13:33:30	ADMIN	Logout.aspx	Logout	Success	N/A
05/05/2010 12:56:24	ADMIN	Logout.aspx	Logout	Success	N/A
05/05/2010 12:48:04	ADMIN	SessionTimeout.aspx	Logout	Success	N/A
05/05/2010 12:40:57	ADMIN	Default.aspx	Login	Success	N/A
05/05/2010 12:40:53	CSR	Logout.aspx	Logout	Success	N/A
05/05/2010 12:39:24	CSR	Default.aspx	Login	Success	N/A
05/05/2010 12:37:33	CSR	Logout.aspx	Logout	Success	N/A
05/05/2010 12:20:20	CSR	Default.aspx	Login	Success	N/A
05/05/2010 12:20:17	ADMIN	Logout.aspx	Logout	Success	N/A
05/05/2010 12:20:01	ADMIN	Default.aspx	Login	Success	N/A
05/05/2010 11:45:47	ADMIN	Default.aspx	Login	Success	N/A
05/05/2010 11:43:30	ADMIN	Default.aspx	Login	Success	N/A
05/05/2010 11:42:11	ADMIN	SessionTimeout.aspx	Logout	Success	N/A
05/05/2010 11:42:11	ADMIN	SessionTimeout.aspx	Logout	Success	N/A
05/05/2010 11:42:11	ADMIN	SessionTimeout.aspx	Logout	Success	N/A
05/05/2010 11:29:34	ADMIN	Default.aspx	Login	Success	N/A
05/05/2010 11:29:24	ADMIN	Logout.aspx	Logout	Success	N/A
05/05/2010 11:28:35	ADMIN	Default.aspx	Login	Success	N/A
05/05/2010 11:27:16	ADMIN	Default.aspx	Login	Success	N/A
05/05/2010 11:27:11	ADMIN	Logout.aspx	Logout	Success	N/A
05/05/2010 11:17:36	ADMIN	Default.aspx	Login	Success	N/A
05/05/2010 10:59:07	ADMIN	Default.aspx	Login	Success	N/A
05/05/2010 10:58:41	ADMIN	Default.aspx	Login	Success	N/A

Figure 40. User Activity Report - Logs

Field	Description
Timestamp	Timestamp for the user event.
User	The user ID associated with the event.
Page	The accessed URL.
Event	The event that occurred on the page.
Return Code	The results of the event.

Report Value
If any values are generated by the audit event, they are inserted into this column.

Activity Counts tab

In the Activity Counts tab, you can review counts of the selected audit events in a time-series chart.

- The selected audit events are listed as individual series data.

Below the chart are displayed break-downs of counts in smaller time periods.

- Reports that span a single day feature count break-downs by the hour. Reports that span multiple days display daily count break-downs.

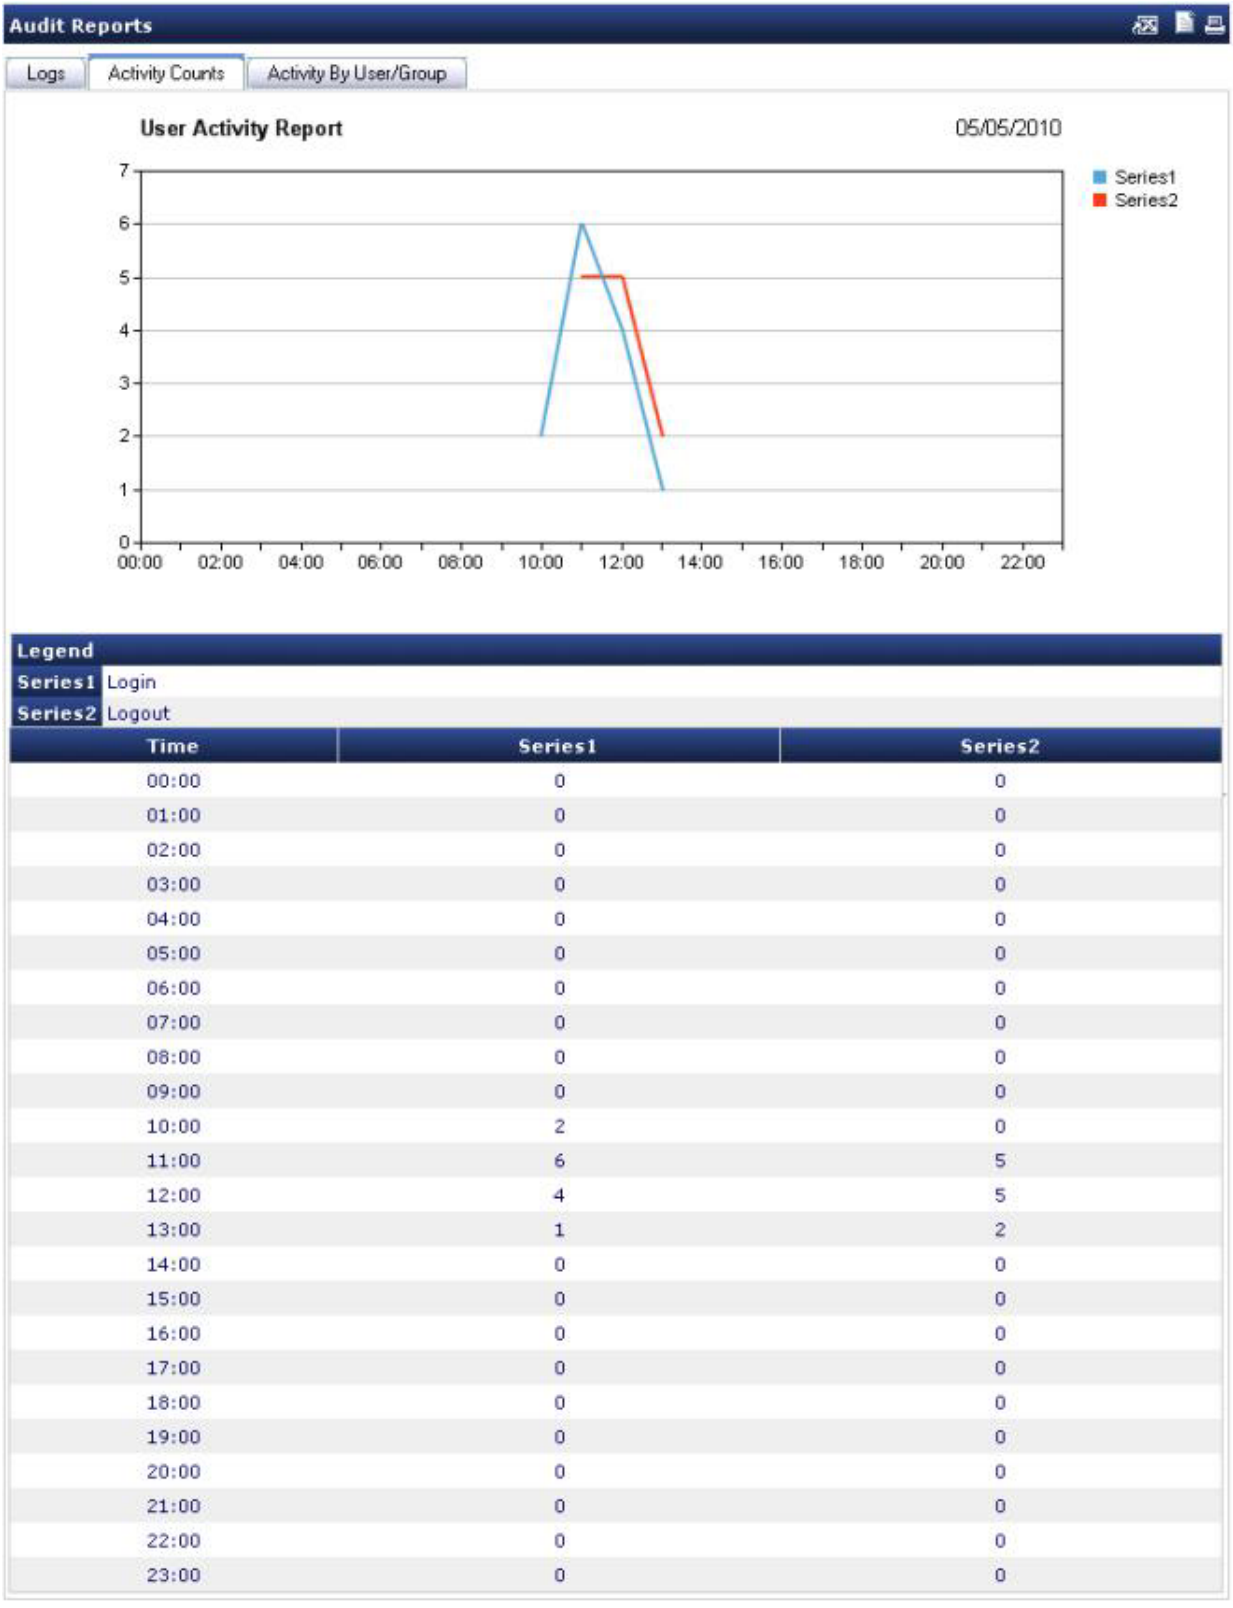

Time	Series1	Series2
00:00	0	0
01:00	0	0
02:00	0	0
03:00	0	0
04:00	0	0
05:00	0	0
06:00	0	0
07:00	0	0
08:00	0	0
09:00	0	0
10:00	2	0
11:00	6	5
12:00	4	5
13:00	1	2
14:00	0	0
15:00	0	0
16:00	0	0
17:00	0	0
18:00	0	0
19:00	0	0
20:00	0	0
21:00	0	0
22:00	0	0
23:00	0	0

Figure 41. User Activity Report - Activity Counts

Activity by User/Group tab

In this tab, you can review total counts of audit events by user. For the selected user, the counts of each selected audit event for the focus period are displayed in individual columns.

- To aggregate user data to the group level, select the Aggregate by User Group checkbox.

  Note: Because individual users may belong to multiple groups, aggregated counts may not match individual user counts.

Figure 42. User Activity Report - Acitivity by User/Group

User Inactivity Report

In the User Inactivity report, you can review Tealeaf user accounts that have not been used during the focus period.

This report is useful for identifying accounts that are no longer in use. For example, using the User Login event, you can identify the users that have failed to login during the selected focus period.

Figure 43. User Inactivity Report

Field Description

Username
The username.

Last Occurrence
For the selected audit event, the timestamp for the last occurrence of the selected audit event or events.

Last Portal Activity
For the selected user, the timestamp of the last recorded activity in the Tealeaf Portal during the focus period.

Usage Metrics Summary Report

This report summaries Tealeaf usage by group for a predefined set of metrics over the focus period.

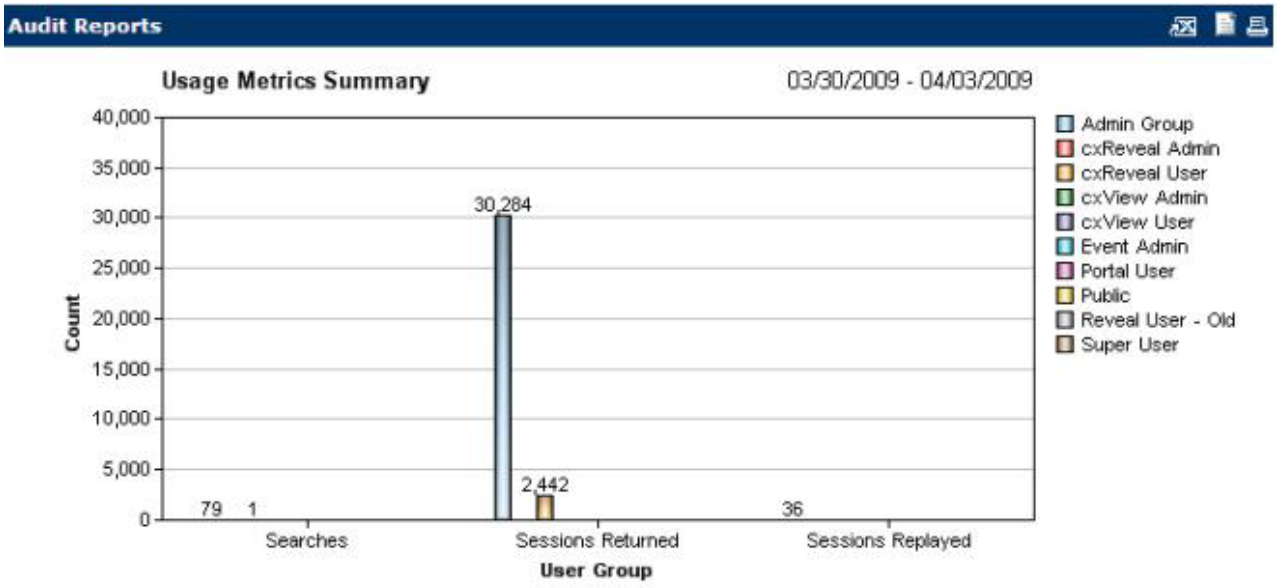

Group	Searches	Sessions Returned	Sessions Replayed	Search Effectiveness %	Replay Use %	Last Login	Last Search	Last Replay
Admin Group	79	30,284	36	38,334.17	45.56	04/03/2009 23:50:52	04/02/2009 23:50:44	N/A
cxReveal Admin	0	0	0	0.00	0.00	03/31/2009 22:07:47	N/A	N/A
cxReveal User	1	2,442	0	244,200.00	0.00	03/31/2009 21:56:06	03/31/2009 21:56:21	N/A
cxView Admin	0	0	0	0.00	0.00	N/A	N/A	N/A
cxView User	0	0	0	0.00	0.00	N/A	N/A	N/A
Event Admin	0	0	0	0.00	0.00	N/A	N/A	N/A
Portal User	0	0	0	0.00	0.00	N/A	N/A	N/A
Public	0	0	0	0.00	0.00	04/01/2009 17:08:50	N/A	N/A
Reveal User - Old	0	0	0	0.00	0.00	N/A	N/A	N/A
Super User	0	0	0	0.00	0.00	N/A	N/A	N/A

Figure 44. Usage Metrics Summary Report

- For more information on the fields in this report, see "Common report fields" on page 62.

Usage Metrics by User Report

The Usage Metrics by user report summarizes usage statistics by individual users over the focus period.

Audit Reports

User	Searches	Sessions Returned	Sessions Replayed	Search Effectiveness %	Replay Use %	Last Login	Last Search	Last Replay
ADMIN	197	91,516	80	46,454.82	40.60	04/07/2009 23:35:04	04/07/2009 23:35:43	N/A
CSR	6	12,406	0	206,766.66	0.00	04/07/2009 19:13:19	04/07/2009 19:13:35	N/A
DEMO	0	0	0	0.00	0.00	N/A	N/A	N/A
KIRK	0	0	0	0.00	0.00	N/A	N/A	N/A
REVEALADMIN	0	0	0	0.00	0.00	03/31/2009 22:07:47	N/A	N/A
TEST	0	0	0	0.00	0.00	04/01/2009 17:08:50	N/A	N/A
VIEW	0	0	0	0.00	0.00	N/A	N/A	N/A
VIEWADMIN	0	0	0	0.00	0.00	N/A	N/A	N/A

Figure 45. Usage Metrics by User Report

- For more information on the fields in this report, see "Common report fields" on page 62.

User Searches Report

This report indicates the searches conducted by user groups over the focus period.

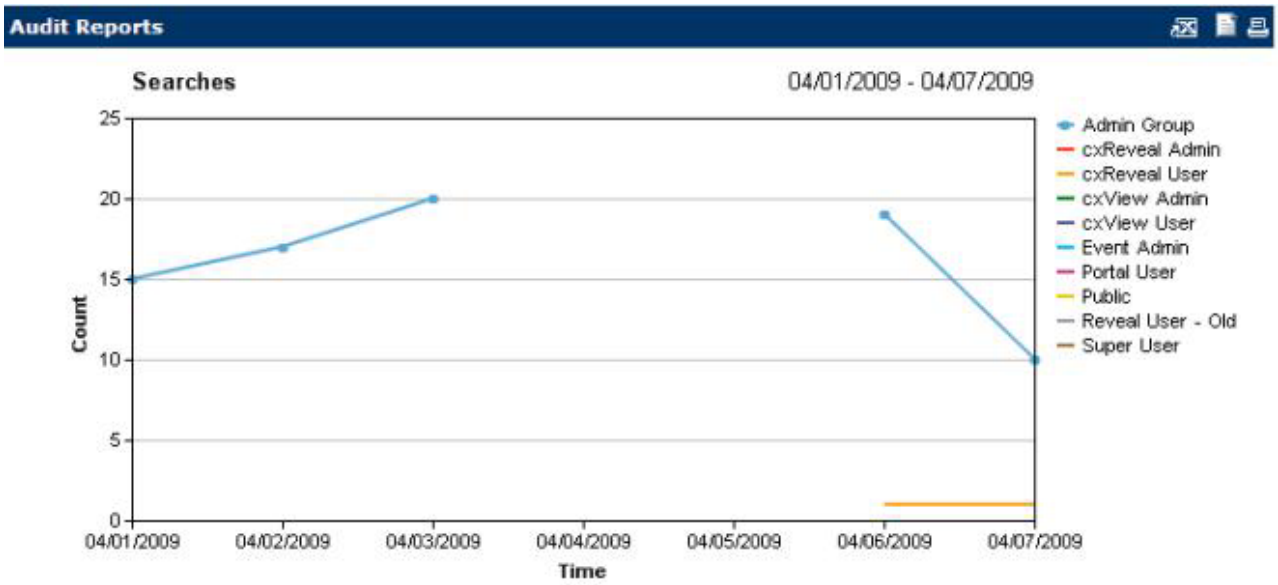

Date	Group	Searches	Sessions Returned	Sessions Replayed	Search Effectiveness %	Replay Use %
04/01/2009	Admin Group	15	6,916	2	46,106.66	13.33
04/01/2009	cxReveal Admin	0	0	0	0.00	0.00
04/01/2009	cxReveal User	0	0	0	0.00	0.00
04/01/2009	cxView Admin	0	0	0	0.00	0.00
04/01/2009	cxView User	0	0	0	0.00	0.00
04/01/2009	Event Admin	0	0	0	0.00	0.00
04/01/2009	Portal User	0	0	0	0.00	0.00
04/01/2009	Public	0	0	0	0.00	0.00
04/01/2009	Reveal User - Old	0	0	0	0.00	0.00
04/01/2009	Super User	0	0	0	0.00	0.00
04/02/2009	Admin Group	17	2,839	1	16,700.00	5.88
04/02/2009	cxReveal Admin	0	0	0	0.00	0.00
04/02/2009	cxReveal User	0	0	0	0.00	0.00

Figure 46. User Searches Report

- For more information on the fields in this report, see "Common report fields" on page 62.
- To review this report for individual users, click the name of the user group. See "Usage Metrics by User Report" on page 69.

Searches with Results Report

This report displays the basic User Searches report with data constrained to only those searches that returned a non-zero number of results.

- See "User Searches Report."

Sessions Returned Report

The Sessions Returned report identifies the number of sessions returned via search by user group over the focus period.

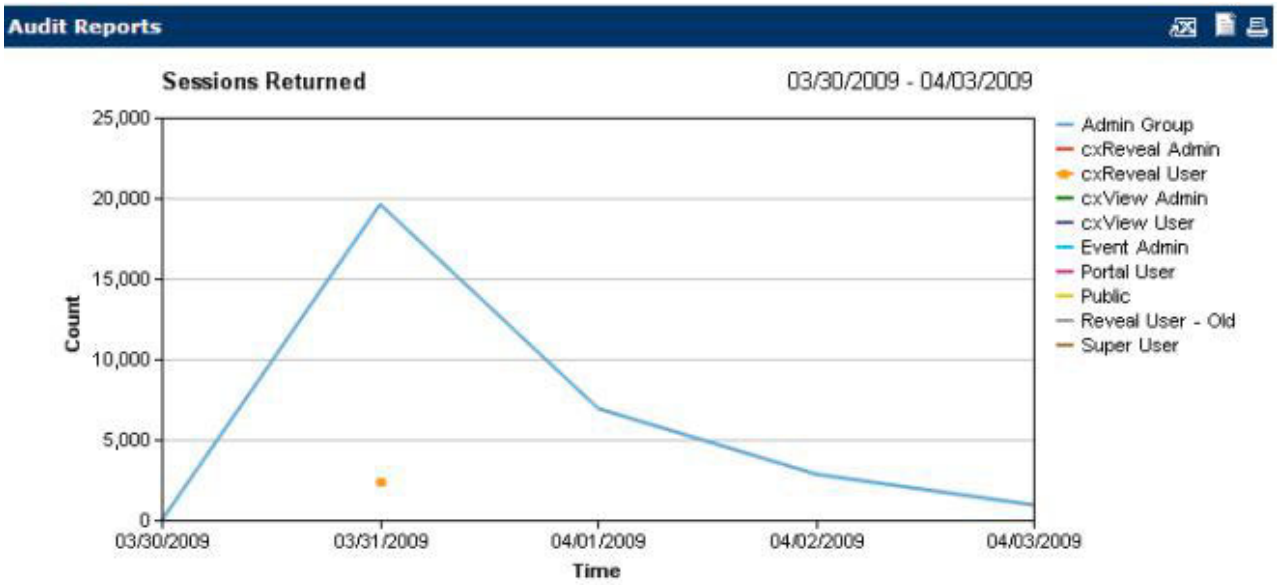

Date	Group	Searches	Sessions Returned	Sessions Replayed	Search Effectiveness %	Replay Use %
03/30/2009	Admin Group	7	12	2	171.42	28.57
03/30/2009	cxReveal Admin	0	0	0	0.00	0.00
03/30/2009	cxReveal User	0	0	0	0.00	0.00
03/30/2009	cxView Admin	0	0	0	0.00	0.00
03/30/2009	cxView User	0	0	0	0.00	0.00
03/30/2009	Event Admin	0	0	0	0.00	0.00
03/30/2009	Portal User	0	0	0	0.00	0.00
03/30/2009	Public	0	0	0	0.00	0.00
03/30/2009	Reveal User - Old	0	0	0	0.00	0.00
03/30/2009	Super User	0	0	0	0.00	0.00
03/31/2009	Admin Group	20	19,594	23	97,970.00	115.00
03/31/2009	cxReveal Admin	0	0	0	0.00	0.00
03/31/2009	cxReveal User	1	2,442	0	244,200.00	0.00
03/31/2009	cxView Admin	0	0	0	0.00	0.00
03/31/2009	cxView User	0	0	0	0.00	0.00
03/31/2009	Event Admin	0	0	0	0.00	0.00
03/31/2009	Portal User	0	0	0	0.00	0.00
03/31/2009	Public	0	0	0	0.00	0.00
03/31/2009	Reveal User - Old	0	0	0	0.00	0.00
03/31/2009	Super User	0	0	0	0.00	0.00

Figure 47. Sessions Returned Report

- For more information on the fields in this report, see “Common report fields” on page 62.
- To review this report for individual users, click the name of the user group. See “Usage Metrics by User Report” on page 69.

Sessions Replayed Report

This report indicates the number of sessions replayed from sessions returned by search for each user group over the focus period.

- This report includes replay requests for both BBR and RTV made from the Portal.
- RTV does not log replay requests inbound from the Portal on its own.

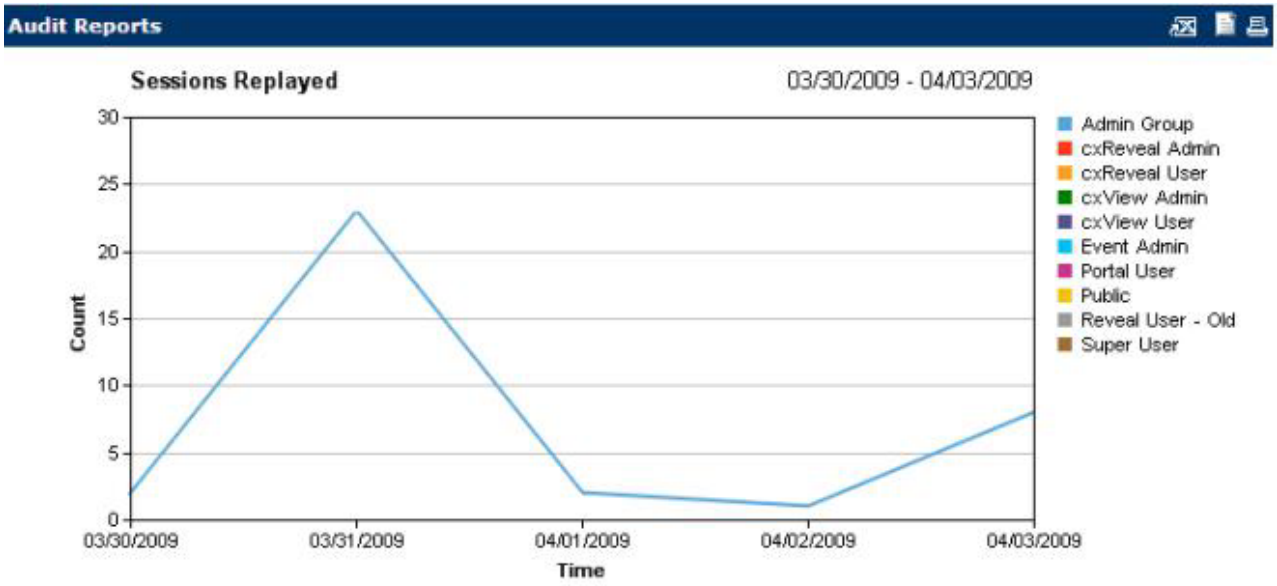

Date	Group	Searches	Sessions Returned	Sessions Replayed	Search Effectiveness %	Replay Use %
03/30/2009	Admin Group	7	12	2	171.42	28.57
03/30/2009	cxReveal Admin	0	0	0	0.00	0.00
03/30/2009	cxReveal User	0	0	0	0.00	0.00
03/30/2009	cxView Admin	0	0	0	0.00	0.00
03/30/2009	cxView User	0	0	0	0.00	0.00
03/30/2009	Event Admin	0	0	0	0.00	0.00
03/30/2009	Portal User	0	0	0	0.00	0.00
03/30/2009	Public	0	0	0	0.00	0.00
03/30/2009	Reveal User - Old	0	0	0	0.00	0.00
03/30/2009	Super User	0	0	0	0.00	0.00
03/31/2009	Admin Group	20	19,594	23	97,970.00	115.00
03/31/2009	cxReveal Admin	0	0	0	0.00	0.00
03/31/2009	cxReveal User	1	2,442	0	244,200.00	0.00
03/31/2009	cxView Admin	0	0	0	0.00	0.00

Figure 48. Sessions Replayed Report

- For more information on the fields in this report, see “Common report fields” on page 62.
- To review this report for individual users, click the name of the user group. See “Usage Metrics by User Report” on page 69.

Searches without Results Report

This report displays the basic User Searches report with data constrained to only those searches that returned zero results.

- See “User Searches Report” on page 71.

Search Effectiveness Report

This metric computes the effectiveness of user-initiated searches. Search effectiveness is computed as number of sessions returned per executed search.

Search Effectiveness can be used to measure how well a customer service representative, for example, is specifying searches. If the CSR has a low Search Effectiveness rating, additional training may help the individual to return searches with a positive number of sessions.

Formula

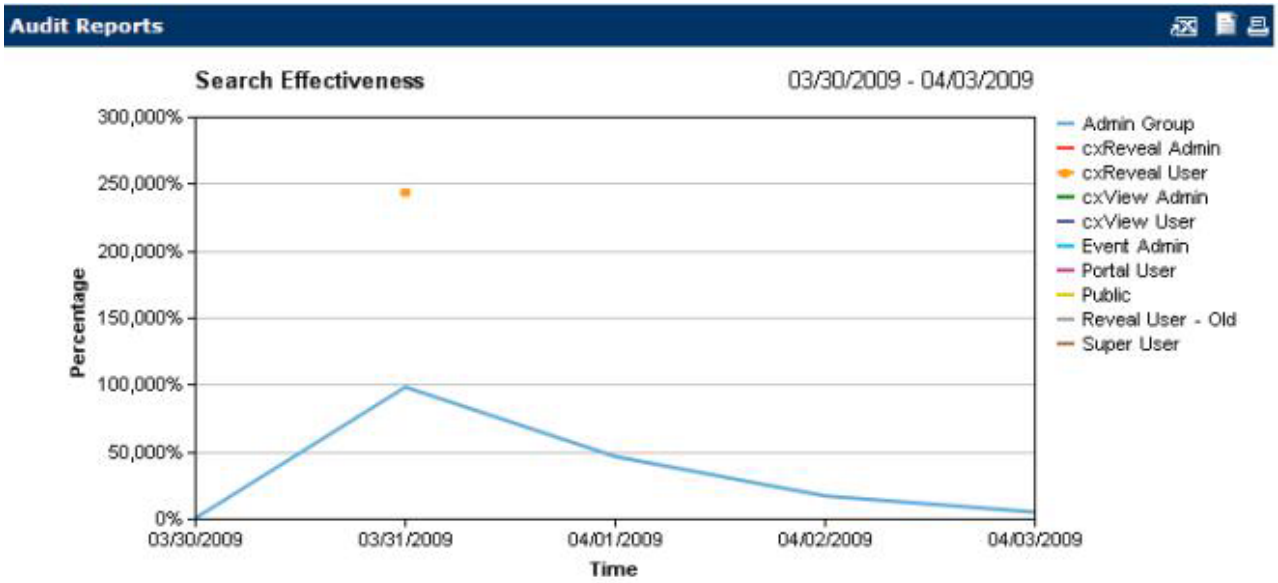

Date	Group	Searches	Sessions Returned	Sessions Replayed	Search Effectiveness %	Replay Use %
03/30/2009	Admin Group	7	12	2	171.42	28.57
03/30/2009	cxReveal Admin	0	0	0	0.00	0.00
03/30/2009	cxReveal User	0	0	0	0.00	0.00
03/30/2009	cxView Admin	0	0	0	0.00	0.00
03/30/2009	cxView User	0	0	0	0.00	0.00
03/30/2009	Event Admin	0	0	0	0.00	0.00
03/30/2009	Portal User	0	0	0	0.00	0.00
03/30/2009	Public	0	0	0	0.00	0.00
03/30/2009	Reveal User - Old	0	0	0	0.00	0.00
03/30/2009	Super User	0	0	0	0.00	0.00
03/31/2009	Admin Group	20	19.594	23	97,970.00	115.00
03/31/2009	cxReveal Admin	0	0	0	0.00	0.00
03/31/2009	cxReveal User	1	2.442	0	244,200.00	0.00

Figure 49. Search Effectiveness Report

- For more information on the fields in this report, see "Common report fields" on page 62.
- To review this report for individual users, click the name of the user group. See "Usage Metrics by User Report" on page 69.

Search Efficiency Report

The Search Efficiency report enables monitoring of how efficient users are in returning manageable number of sessions in each search. Efficiency is calculated as the inverse of the average number of sessions returned per search.

A smaller number of returned sessions yields a higher efficiency rating in finding sessions, which is a useful metric particularly for customer service representatives.

For example, if a CSR is able to execute a search that returns one (1) session, then that individual has been efficient in searching. However, if that individual subsequently returns 100 sessions per search, his or her efficiency rating should drop, which may correlate to a longer time to resolve the customer issue and an impact on customer satisfaction.

Formula

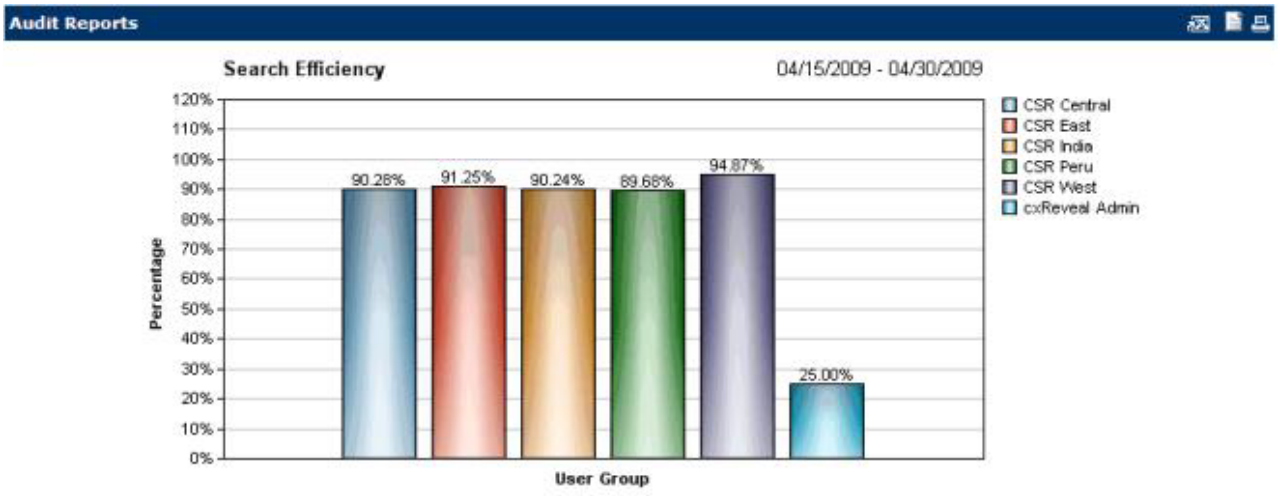

Group	Searches	Searches with Results	Sessions Returned	Sessions Replayed	Search Effectiveness %	Search Efficiency %	Replay Use %	Searches without Results	Last Login	Last Search	Last Replay
CSR Central	76	65	72	54	85.52	90.27	83.07	11	05/28/2009 20:44:27	05/28/2009 18:47:31	05/28/2009 20:46:13
CSR East	174	146	160	151	83.90	91.25	103.42	28	05/28/2009 23:49:33	05/28/2009 18:58:19	05/29/2009 05:41:28
CSR India	44	37	41	39	84.09	90.24	105.40	7	05/28/2009 17:54:47	N/A	05/28/2009 21:26:45
CSR Peru	117	113	126	123	96.58	89.68	108.84	4	05/29/2009 01:00:18	05/28/2009 22:52:15	05/28/2009 23:46:24
CSR West	84	37	39	32	44.04	94.87	86.48	47	05/29/2009 00:04:19	N/A	05/29/2009 04:18:44
cxReveal Admin	1	1	4	1	100.00	25.00	100.00	0	04/29/2009 00:11:33	04/29/2009 00:11:41	04/29/2009 00:11:45

Figure 50. Search Efficiency Report

- For more information on the fields in this report, see "Common report fields" on page 62.
- To review this report for individual users, click the name of the user group. See "Usage Metrics by User Report" on page 69.

Replay Use Report

This report displays the number of replays executed by each user group over the focus period.

Formula

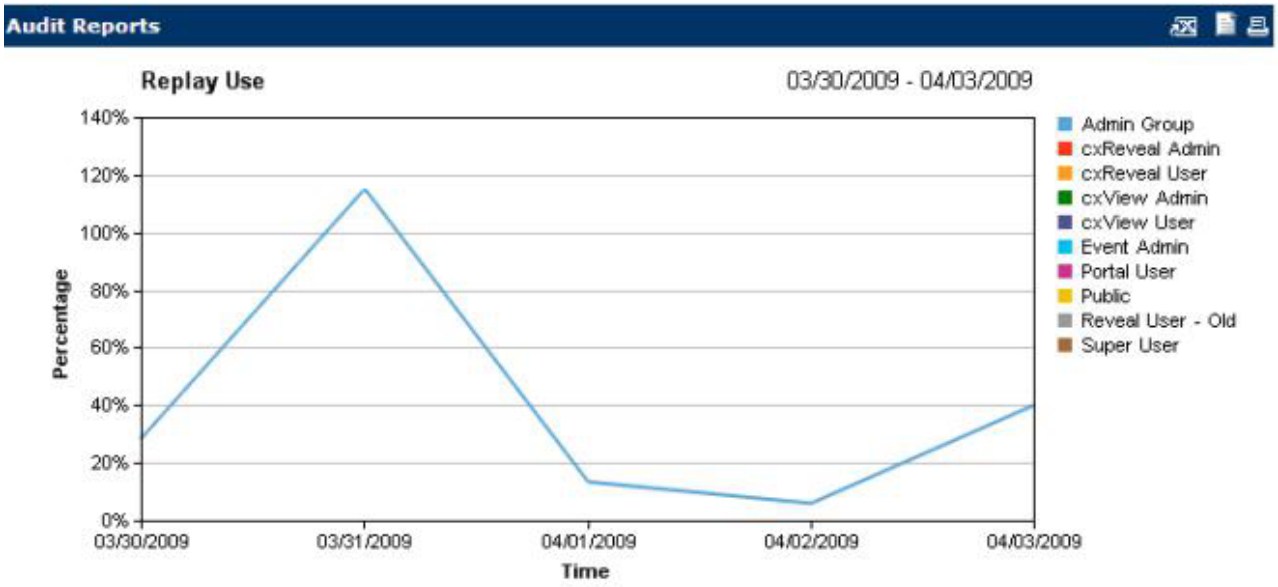

Date	Group	Searches	Sessions Returned	Sessions Replayed	Search Effectiveness %	Replay Use %
03/30/2009	Admin Group	7	12	2	171.42	28.57
03/30/2009	cxReveal Admin	0	0	0	0.00	0.00
03/30/2009	cxReveal User	0	0	0	0.00	0.00
03/30/2009	cxView Admin	0	0	0	0.00	0.00
03/30/2009	cxView User	0	0	0	0.00	0.00
03/30/2009	Event Admin	0	0	0	0.00	0.00
03/30/2009	Portal User	0	0	0	0.00	0.00
03/30/2009	Public	0	0	0	0.00	0.00
03/30/2009	Reveal User - Old	0	0	0	0.00	0.00
03/30/2009	Super User	0	0	0	0.00	0.00
03/31/2009	Admin Group	20	19,594	23	97,970.00	115.00
03/31/2009	cxReveal Admin	0	0	0	0.00	0.00
03/31/2009	cxReveal User	1	2,442	0	244,200.00	0.00

Figure 51. Replay Use Report

- For more information on the fields in this report, see "Common report fields" on page 62.
- To review this report for individual users, click the name of the user group. See "Usage Metrics by User Report" on page 69.

Export

In the toolbar, click the Export tool.

- **Excel** - Export the report data to a single Excel worksheet in XML format.
- **PDF** - Export the report to PDF format. Links in the generated report point back to the Tealeaf Portal page.
- **Email** - Generate the report in either format and send it via email:

1. Select the report format: PDF or Excel.
2. Enter an optional message.
3. Enter recipient email addresses or aliases, separated by commas.
4. To send the report, click **Send**. The report is exported in the selected format and emailed to the specified list of addresses.

Portal Application Information Report

Through the Portal, you can review basic configuration information about the Portal application and execute a series of tests of various components of the Portal application.

This report is available to all Tealeaf users based on the type of user: administrator or basic user report. Below, the contents of each pane are listed for the administrator version, which contains all available Portal application information.

- This report requires no configuration.
- To review the Portal Application Information report, select **Help** > **About IBM Tealeaf CX Portal**.

Portal Performance Tests

In the Portal Performance Tests panel, you can review a set of performance tests that can be applied to the Tealeaf Portal application. These tests can be used to identify Portal application performance in various browser configurations and network scenarios.

- To execute any of the following tests, click the Execute link next to the test name.
- To execute all available tests, click the Execute All link.

Test Name
Description

AJAX Toolkit Test 1
Tests ability of using AJAXPro toolkit, which is used by the Portal application, in your web browser. Results are `Success` or `Failure`.

AJAX Toolkit Test 2
Tests performance of ComponentArt toolkit, which is used by the Portal application, in your web browser. Results are `Success` or `Failure`.

Data Transfer (1 Byte)
Tests the time to transfer one byte between your web browser and the web server and back again. Results are reported in milliseconds.

Data Transfer (100 KB)
Tests the time to transfer 100 kilobytes between your web browser and the web server and back again. Results are reported in milliseconds.

Data Transfer (1 MB)
Tests the time to transfer 1 megabyte between your web browser and the web server and back again. Results are reported in milliseconds.

Simple Database Query
Performs a simple database query of the Tealeaf Reports database. Results are reported in milliseconds.

- A value of `0.00 ms` indicates that the Reports database is hosted on the same machine as the Portal application.

Simple JavaScript Benchmark
Performs a JavaScript benchmark test of a simple JavaScript in your web browser, indicating the time in milliseconds to execute the JavaScript.

Complex JavaScript Benchmark
Performs a JavaScript benchmark test of a more complex JavaScript in your web browser, indicating the time in milliseconds to execute the JavaScript.

Database Information

This section identifies all of the Tealeaf databases to which the Portal application is connected. For each database, you can review the Tealeaf build version, the installation date and time, and the current size of the database.

Database Key
Database Name

TL_SYSTEM
Tealeaf system database

TL_REPORTS
Tealeaf Reports database

TL_RSEXTRACTOR
Tealeaf Result Set Extractor staging database

Note: This database only exists in new installations if you have licensed IBM Tealeaf cxResults. In upgraded systems, this database may already exist.

TL_SEARCH
IBM Tealeaf cxReveal search database(s). There may be multiple `TL_SEARCH` databases.

Note: This database only exists if you have licensed IBM Tealeaf cxReveal and have installed the IBM Tealeaf cxReveal Server and database. See "cxReveal Installation" in the *IBM Tealeaf cxReveal Administration Manual*.

TL_STATISTICS
Tealeaf Statistics database

TL_VISSTAGE
Tealeaf Visitor Staging database

Note: The Visitor database is available only if you have licensed IBM Tealeaf cxResults. please contact your IBM Tealeaf representative.

TL_VISREPORT
Tealeaf Visitor Reporting database

Note: The Visitor database is available only if you have licensed IBM Tealeaf cxResults. please contact your IBM Tealeaf representative.

Portal Component Names

This panel lists the names and versions of the components included in the Tealeaf Portal application. This section is useful for debugging purposes only.

Tealeaf Administrator

The Tealeaf administrator's contact information is listed on the page. For questions or comments regarding the Tealeaf system, contact your Tealeaf administrator.

- This information can be configured from within the Portal. For more information, see the Miscellaneous settings category in the “CX Settings” on page 192 page.

Time Zone Offset

This panel indicates the time zone offsets from Greenwich Mean Time (GMT) to the Tealeaf system time.

Time zone offsets are listed as + or - GMT time (e.g. GMT-7), and current time stamps are listed.

- See "Configuring the System Timezone" in the *IBM Tealeaf CX Configuration Manual*.

System Properties

You can review various system properties for the host of the Portal Server.

Property
Description

OS Name
Operating System name

OS Service Pack
If a Service Pack has been applied to the operating system installed on the Portal server, the service pack name is listed.

OS Version
Version information for the operating system

OS Architecture
Architecture information for the operating system. 64-bit versions are listed as x64.

OS UI Culture
Locale setting in use by the operating system

OS UI Culture NumberFormat
Method for formatting number groups and the decimal point by the operating system

Thread Culture
Locale setting in use when parsing input data

Thread Culture NumberFormat
Method for formatting number groups and the decimal point when parsing input data

Thread UI Culture
Locale setting in use when formatting data for display

Thread UI Culture NumberFormat
Method for formatting number groups and the decimal point when formatting data for display

Virtual Machine Host
If the Portal Server is installed in a virtual machine, the VM host is listed here.

Note: Installing Tealeaf software in a VM is not a recommended configuration. See "CX Pre-Installation Checklist" in the *IBM Tealeaf CX Installation Manual*.

Total Memory
Total RAM on the Portal Server host

Avail Memory
RAM available to the Portal application on the host

Host Name
The name of the host

Fully Qualified Domain Name (FQDN)
Fully qualified domain name for the Portal application

IPs Internet Protocol address for the host

NICx For each Network Interface Card connected to the Portal Server, you can review name, description, and type information, as well as status and current transfer rate.

CPUx For each CPU installed on the Portal Server machine, you can review the processor type and performance rating, as well as current performance rating.

Drive: x:
Information for the storage drive where the Tealeaf Portal application is installed.

UI Culture
Language configured for the Portal application.

ASP.NET Version
Current version of ASP.NET in use by the Portal application.

ASP.NET Cache Limit
Limit to the size of an individual ASP.NET cache.

ASP.NET Cache Count
Number of cached files currently in use by the installed Portal application.

Monitoring Tealeaf Using External Tools

Several external, publicly available tools can be used to monitor Tealeaf services and overall performance, as well as Tealeaf database activities.

Basic Connectivity

- "Web browser"
- "Nslookup" on page 82
- "Tracert" on page 82

Web browser

You can test the availability of connections within the enterprise network infrastructure by inserting a Tealeaf server and port number as the address in your web browser.

For example, to test the availability of the PCA Web Console, enter the following in your address bar:

```
https://<PCA_server>:8443
```

where: `<PCA_server>` is the hostname of the PCA server.

Nslookup

`Nslookup.exe` is a Microsoft command-line tool that can be used to troubleshoot network connectivity and DNS issues.

- For more information, see `http://support.microsoft.com/kb/200525`.

Tracert

Traceroute can be used to diagnose the internet gateway hops between a selected server and a specified IP address.

- For more information, see `http://www.tracert.com`.

Capture

- "tcpdump"
- "Wireshark"

tcpdump

Tcpdump captures network packets and optionally dumps them to a file for offline analysis. This tool is useful for detecting and resolving network connectivity issues and verifying data processed by the PCA.

- For more information, see `http://www.tcpdump.org`.

Wireshark

Wireshark can be used to read and write a variety of capture formats from Windows or Linux operating systems. This tool provides a GUI with which you can review and analyze network traffic captured by Tcpdump.

- For more information, see `http://www.wireshark.org`.

Windows Servers

- "PerfMon"
- "Process Explorer" on page 83
- "Windows Services Control Panel"

PerfMon

The Windows Performance Monitor (PerfMon) can be used to monitor performance of individual Tealeaf services on a single-server or multi-core server.

- To launch PerfMon, from the Start menu, select **Run...**. Enter `perfmon` and press `ENTER`.
- For more information, see the Help menu in PerfMon.

Windows Services Control Panel

On any Windows server hosting Tealeaf software, you can review the status of all Tealeaf services through this Control Panel.

Note: TMS is the preferred way to stop, start, and restart all Tealeaf components. See Chapter 4, "Tealeaf Management System," on page 259.

Process Explorer

This add-on utility enables you to explore the processes that have been opened by Tealeaf DLLs. Process Explorer is useful for identifying problems with DLL performance, versioning, and other issues

- When it is installed, you can access Process Explorer from the Windows desktop. Click the desktop and press `CTRL + SHIFT + ESC`.
- For more information, see `http://technet.microsoft.com/en-us/sysinternals/bb896653.aspx`.

Database

- "SQL Manager"
- "SQLIO"

SQL Manager

SQL Manager for SQL Server provides a visual interface for developing and deploying databases and includes useful wizards for common tasks.

- For more information, see `http://sqlmanager.net/en/products/mssql/manager`.

SQLIO

Microsoft provides a useful tool for monitoring input and output of SQL Server databases.

- For more information, search `http://www.microsoft.com` for "SQLIO Disk Subsystem Benchmark Tool".
- For command line documentation of SQLIO, see `http://technet.microsoft.com/en-us/library/cc966412.aspx`.

Replay

- "Fiddler"
- "YSlow" for Firefox
- "HTTPwatch" for Internet Explorer and Firefox

Fiddler

Fiddler is a debugging tool for analyzing all HTTP or HTTPS traffic between your browser and the Internet.

- For more information, see `http://www.fiddler2.com/fiddler2`.

YSlow

The Yahoo! YSlow tool analyzes web pages and suggests methods for improving their performance. This Firefox add-in is integrated with the Firebug web development tool.

- For more information, see `http://developer.yahoo.com/yslow`.

HTTPwatch

This utility enables viewing and debugging of HTTP traffic and is integrated with Internet Explorer and Firefox.

- For more information, see `http://www.httpwatch.com`.

Reporting for Tealeaf Administrators

This section provides an overview of resources available to Tealeaf system administrators who wish to monitor overall system health and individual Tealeaf components. When an issue arises, the links in this section can assist you in quickly assessing the issue and beginning the troubleshooting process.

Monitoring Current System Status

This section provides references to areas of the Tealeaf Portal through which you can monitor the current status of various aspects of the Tealeaf system.

- For more information on how to use these reports, see "Troubleshooting Tealeaf CX Issues through the Portal" on page 97.
- For Tealeaf best practices on monitoring application health, see https://community.tealeaf.com/display/res/Monitoring+Application+Health.

Active menu

In the Portal, the Active menu provides access to a number of monitoring tools that can be used to assess how data is flowing through the system.

- `Status` - Review current status of each Tealeaf server or all servers with respect to session and error counts, page generation times, and data volume. See "Monitoring Active Status" in the *IBM Tealeaf cxImpact User Manual*.
- `Sessions` - Drill into the currently active sessions of the monitored application. You can review individual pages and replay the session. See "Monitoring Active Sessions" in the *IBM Tealeaf cxImpact User Manual*.
- `Events` - Review the events that have been encountered in the active traffic stream within the last 30 minutes. See "Monitoring Active Events" in the *IBM Tealeaf cxImpact User Manual*.
  - Events are created through the Tealeaf Event Manager. See "TEM Events Tab" in the *IBM Tealeaf Event Manager Manual*.

Through the Alert Monitor, you can review the current status of all enabled alerts, including the most recently generated instance of each one. See "Alert Monitor" in the *IBM Tealeaf cxImpact User Manual*.

- Alerts are created through the Tealeaf Event Manager. See "TEM Alerts Tab" in the *IBM Tealeaf Event Manager Manual*.

System Dashboards

By default, Tealeaf includes the following useful dashboards to provide a top-down perspective on system status and metrics:

These dashboards are available through the Dashboards menu in the Portal for IBM Tealeaf cxView users and in the menus listed below for IBM Tealeaf cxImpact users.

Note: Access to Dashboards requires IBM Tealeaf cxView, a separately licensable component of the IBM Tealeaf CX system. please contact your IBM Tealeaf representative.

1. **Technical Site Metrics** - This report provides useful information on system and user activities, as well as reports on individual pages. See "Technical Site Metrics Dashboard" in the *IBM Tealeaf Reporting Guide*.

2. **Technical System Status** - This dashboard summarizes technical status for current sessions, Tealeaf canisters, the DecoupleEx session agent, storage usage, and database size.
   - You can also access this dashboard by selecting **Tealeaf** > **System Status** > **Status Summary**. See “System Status Report.”
   - In the System Status menu, you can access reports on individual components of the Tealeaf system. See “System Status Report.”

For more information on dashboards, see "Using Dashboards" in the *IBM Tealeaf cxView User Manual*. IBM Tealeaf cxView administrators can configure dashboards to meet their needs. See "Configuring Dashboards" in the *IBM Tealeaf cxView User Manual*.

Activity Reports

Tealeaf provides a set of activity reports that are accessible through the Tealeaf Report Builder. These reports provide detailed metrics on the data that Tealeaf has captured from your web application.

- See "Reports Provided by Tealeaf" in the *IBM Tealeaf Reporting Guide*.

Performance Reports

Performance reports provide insight into how your web application is performing with respect to content delivery to your visitors.

- See "Reports Provided by Tealeaf" in the *IBM Tealeaf Reporting Guide*.

System Status Report

In the System Status report, you can review the current system status for the following components of the Tealeaf system for the selected server:

- **Canister** - Current canister information includes counts of captured sessions and events and sessions awaiting evaluation, indexing, and archiving. You can also review canister and memory usage statistics.
- **DecoupleEx** - Current processing status of the Extended Decoupler Session Agent, if in use on the server. Information includes hits and megabytes of data being passed into the canister and out of the canister.
- **Storage** - The Storage report indicates storage usage for any canisters or indexes maintained on the server.
- **Health Based Routing** - If HBR is enabled on the server, this report indicates HBR distribution patterns and the current status information of each downstream Processing Server.
- **Database Filegroup Size** - Provides report on size and growth of database filegroups for each database.
- **Database Table Size** - Reports on size and growth of each table in type of data group.
- **Collection Status** - This separate report details the results of the most recent runs executed by the Data Collector, which gathers data from the available Canisters, aggregates them, and then inserts them into the reporting databases. See “Data Collector Status” on page 94.

See “System Status” on page 8.

These reports are very useful for identifying problems in processing data in the Tealeaf system. See “Troubleshooting Tealeaf CX Issues through the Portal” on page 97.

Checking Component Status through TMS

TMS can be an effective tool for checking which components are currently running.

In the TMS WorldView tab, select `Servers` in the View drop-down. Click the + next to the server you wish to check. The components for the selected server are displayed.

- If a component is running, a green icon is displayed next to the name of the component, and the Stop button appears in the Actions panel.
- If the component is stopped, a gray icon is displayed next to the name of the component. If the component can be started in TMS, a Start button appears in the Actions panel.

Through TMS, you can initiate tasks for specific servers.

- To start, stop, or restart a server, select it in TMS WorldView and then click the appropriate button in the Component Actions pane.
- You can also configure other kinds of jobs through TMS. See Chapter 5, "TMS Jobs Tab," on page 291.

Tracking System Statistics

In addition to the web application session data, Tealeaf can capture hits containing statistics reported from multiple components of the system.

The following components can be configured to report statistics hits for processing and reporting:

- **DecoupleEx** - monitors data that is passed to the Extended Decoupler Session Agent at the beginning of the main processing pipeline.
- **IBM Tealeaf CX Capture** - monitors data captured by the IBM Tealeaf CX Passive Capture Application.
- **CanisterAS** - Monitors activity status on the Short-Term and Long-Term Canisters through a pre-configured set of canister events.

Through the Portal, you can configure and display reports on the above captured data.

- See "System Statistics" on page 39.

Reporting on Tealeaf Servers

Through the Portal, you can access a variety of reporting information on the various servers in the IBM Tealeaf CX environment.

Portal Management Logs

For some servers, you can access system functions and logs through the Portal Management page.

- To access the Portal Management page, select **Tealeaf** > **Portal Management**. To manage your Tealeaf servers, click the Manage Servers link in the left navigation panel.

When you select a Tealeaf server, a toolbar of available tools is displayed at the top of the page. The following tools may be available in the toolbar above the server information.

Table 3. Portal Management Logs

Icon	Log	Description
	Windows Event Log	This tool displays the Windows Event log from the selected server and enables you to filter the display by message type, severity, and time period.
	Tealeaf Logs	Use this tool to see all Tealeaf log files currently stored on the server.

See "Managing Tealeaf Servers" on page 115.

Exporting Tealeaf logs

As needed, you can export selected logs from the Tealeaf system through the Portal for offline analysis.

• See "Portal Logs" on page 227.

Via TMS

Through the Tealeaf Management System, you can acquire additional information pertaining to individual servers.

1. Login to the Portal as an administrator.
2. From the Portal menu, select **Tealeaf** > **TMS**.
3. In the WorldView tab, select the Servers view.
4. Select each server.
5. Server information including whether the server is the TMS master server and the TMS components enabled on the server is displayed in the Server Info panel.
   • See "TMS WorldView Tab" on page 262.

To export the configuration summary of the selected server, click the Advanced tab. Then, click **Export Config Summary**.

• See Chapter 7, "TMS Advanced Tab," on page 303.

Reporting on Tealeaf Services

The Portal Management page contains useful information on the individual services deployed by Tealeaf on each server.

1. To access, select **Tealeaf** > **Portal Management** in the Portal menu.
2. Click the Manage Servers link in the left navigation panel.
3. Select a server whose services you wish to review.
4. **Checking if services are running:** When the server is selected, click the Registry and Services tool ( ). The list of Windows services for the selected server is displayed. A green light icon indicates that the service is enabled and working.
   • See "Managing Tealeaf Servers" on page 115.
5. **Reviewing service logs:** Through the Portal Management page, you can review the log files of various Tealeaf services.

- Click the Logs panel. Click a link to display log information for one of the displayed services.
- See "Portal Logs" on page 227.

Reporting on TMS

The Tealeaf Management System provides centralized management of Tealeaf configurations and can be used to configure and apply tasks to individual servers.

- See Chapter 4, "Tealeaf Management System," on page 259.

Through TMS, you can monitor the running state of each component or service in the system. See "Checking Component Status through TMS" on page 86.

TMS Logs

Logs for the Tealeaf Management System can be retrieved through the Portal Management page.

1. Login to the Portal as an administrator.
2. From the Portal menu, select **Tealeaf** > **Portal Management**.
3. Click the Logs link in the left navigation panel.
4. Click the Tealeaf Logs and Configuration Files link.
5. TMS logs are stored on the Report Server. Download logs from this server to retrieve the TMS log files for the selected time period.
   - See "Portal Logs" on page 227.

TMS in Tealeaf Status

The Tealeaf Status report includes status information from the Tealeaf Management System.

- See "Tealeaf Status Report" on page 89.

Reporting on the Portal Application

In addition to log files, Tealeaf provides the following useful features for reporting on the status of the Portal application itself.

1. **Portal Application Information Report** - Through the Portal, you can access current status and configuration information for the Portal. See "Portal Application Information Report."
2. * Tealeaf Status Report* - The Tealeaf Status report contains statistical information on current states of the active Tealeaf components. See "Tealeaf Status Report" on page 89.
3. **Portal User Activities** - Through the Portal, you can configure and display reports on a wide variety of user activities in the Portal. See "Reporting on Users and Groups" on page 89.

Portal Application Information Report

Through the Portal, you can access report information and testing tools for the Portal application.

- From the Portal menu, select **Help** > **About IBM Tealeaf CX Portal**.
- See "Portal Application Information Report" on page 78.

Tealeaf Status Report

The Tealeaf Status report assembles a variety of status information on Tealeaf servers and components into a single report that is accessible through the Tealeaf Portal.

Steps: Tealeaf administrators can access the Tealeaf Status report through the Tealeaf Portal.

1. Login to the Tealeaf Portal as an administrator.
2. Select **Tealeaf** > **Portal Management**.
3. Click the Logs link.
4. Click **View Tealeaf Status**.
5. The Tealeaf Status report is generated based on the current state of all active servers defined in the Portal Management display. This report is displayed in a new browser window.
   - See "Tealeaf Status Report" on page 231.

Reporting on Users and Groups

This section describes monitoring users and groups.

User Activity reports

Through the Tealeaf Portal, you can monitor specific activities of individual users and groups, including logins, searches, replays, and more.

- From the Portal menu, select **Tealeaf** > **User Activity**.

Report Name
Description

"Monitoring User Activity" on page 57
For the selected user or group, this report displays logs and activity counts.

"Monitoring User Activity" on page 57
Review the Tealeaf accounts that have not been used for the selected focus period.

"Monitoring User Activity" on page 57
Usage by group for a predefined set of metrics over the focus period

"Monitoring User Activity" on page 57
Usage by individual user for a predefined set of metrics over the focus period

"Monitoring User Activity" on page 57
Identifies searches by user group over the focus period. You can drill into individual searches through this report.

"Monitoring User Activity" on page 57
The Searches report constrained to display only searches that returned a non-zero number of results. You can drill into individual searches through this report.

"Monitoring User Activity" on page 57
Count of sessions returned by user group for the focus period

"Monitoring User Activity" on page 57
Count of sessions replayed by user group for the focus period

"Monitoring User Activity" on page 57
The Searches report constrained to display only searches that returned zero results. You can drill into individual searches through this report.

"Monitoring User Activity" on page 57
Identifies the number of returned sessions per executed search.

"Monitoring User Activity" on page 57
The inverse of the average number of sessions returned per user group for the focus period

"Monitoring User Activity" on page 57
Number of replays executed by user group over the focus period

For more information on these reports and how to configure them, see "Monitoring User Activity" on page 57.

User administration

Through the Portal, you can access user and group administration tools, depending on the licensed Tealeaf product.

- **IBM Tealeaf cxImpact:** See "CX User Administration" on page 165.
- **IBM Tealeaf cxView:** See "cxView User Administration" on page 184.
- **IBM Tealeaf cxReveal:** See "cxReveal User Administration" in the *IBM Tealeaf cxReveal Administration Manual*.

Reporting on Captured Hits

A request and its matching response are bundled together as a hit in the IBM Tealeaf CX Passive Capture Application. After the PCA captures a hit, it is forwarded to the Transport Service or the Health-Based Routing machine for disposition to the appropriate Short Term Canister.

After the hit is stored in Canister memory, it is passed through each of the sequential session agents in the pipeline until the final or terminal session agent completes the processing.

- The Tealeaf pipeline is a configurable set of Windows-based pipelines composed of a sequence of agents, each of which performs data manipulation and processing of hits. See "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

This section provides links to useful reporting information on activities before, during, and after pipeline processing.

Hits in the PCA

Through the PCA Web Console, you can monitor statistical information on the capture process and identify potential problems.

- See "PCA Administration Reporting" on page 96.

Hits in the Transport Service or HBR

After a request has been matched with the corresponding response and passed through the PCA's internal pipeline, the PCA forwards the bundled hit to the Windows servers for processing.

Depending on your Tealeaf deployment, the receiver can be either of the following:

- **Tealeaf Transport Service** - This service receives hits from the PCA and manages their processing through the Windows pipeline.
  - See “Monitoring the Pipeline.”
- **Health-Based Router** - In multi-Canister deployments, HBR can be deployed to manage spooling and load balancing between Canisters. See "Health-Based Routing (HBR) Session Agent" in the *IBM Tealeaf CX Configuration Manual*.
  - You can monitor the status of HBR through the System Status page in the Tealeaf Portal. See “System Status” on page 8.

In the Windows Pipeline

This section describes how the Windows pipeline is monitored.

DecoupleEx

At the top of the Windows pipeline, the Extended Decoupler session agent manages the rate at which hits are fed into the pipeline.

- Through the Tealeaf Portal, you can monitor the status of the session agent through the System Status page. See “System Status” on page 8.
- The Extended Decoupler session agent can be configured to send statistics hits for insertion into the Tealeaf database, where they are aggregated for reporting. See “System Statistics” on page 39.
- See "Extended Decoupler Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

Monitoring the Pipeline

You can monitor the activities of the Windows pipeline and the aggregated hit statistics of content being passed through individual session agents through TMS.

In the Pipeline Status tab, you can review current statistics on hit counts and data size as content is passed through each session agent.

- See Chapter 6, “TMS Pipeline Status Tab,” on page 297.
- For more information on how to use the Pipeline Status tab to troubleshoot issues, see “Troubleshooting Tealeaf CX Issues through the Portal” on page 97.

After Pipeline in STC

After a hit has been successfully processed through the Windows pipeline, it is retained in the Short Term Canister, an in-memory database where hits from active sessions are accumulated until the session is closed.

- See “Troubleshooting Tealeaf CX Issues through the Portal” on page 97.
- You can monitor the status of individual canisters through the System Status page. See “System Status” on page 8.

Reporting on LTC

When a session is closed by user exit, session timeout setting, or Tealeaf event, the session is moved from the STC to the Long Term Canister (LTC), which is a disk-based database of sessions.

- Sessions in the LTC are marked for indexing, which is executed as a separate process at predefined intervals. Indexing can be monitored through the Canister Status report in the System Status page. See “System Status” on page 8.
- For more information on indexing, see "Configuring CX Indexing" in the *IBM Tealeaf CX Configuration Manual*.

- For more information on how to troubleshoot issues in indexing, see “Troubleshooting Tealeaf CX Issues through the Portal” on page 97.

Sessions are stored in the Long Term Canister for a predefined number of days. Depending on data volumes of your site, you may need to monitor the current day's storage volume and anticipated storage requirements for the time period in which sessions are stored in the LTC.

- You can monitor LTC storage volumes through the Storage report in the System Status page. See “System Status” on page 8.
- For more information on troubleshooting issues in the LTC, see “Troubleshooting Tealeaf CX Issues through the Portal” on page 97.

Sessions in the LTC can be searched through the Portal and RTV. Reporting on search effectiveness is available for user groups. See “Reporting on Users and Groups” on page 89.

- IBM Tealeaf cxReveal users can review dashboard charts on search effectiveness. See "cxReveal Overview Dashboard" in the *IBM Tealeaf cxReveal User Manual*.

For more information on searching through the Portal, see "Searching Session Data" in the *IBM Tealeaf cxImpact User Manual*.

- IBM Tealeaf cxResults users can search by visitor for session data and other visitor-based information. See "Searching for Visitors" in the *IBM Tealeaf cxResults User Manual*.
- For more information on searching through RTV, see "RealiTea Viewer - Session Search and Subsearch" in the *IBM Tealeaf RealiTea Viewer User Manual*.

Reporting on Tealeaf Databases

Through the Portal, you can review a wide range of reports that are sourced from the Tealeaf databases.

- For more information on these reports, see “Database-Sourced Portal Reports.”
- The Portal provides quite a bit of useful information on database status and the Portal application's interactions with it. See “Portal Database Queries” on page 93.

After sessions have been stored and indexed in the LTC, data from them is collected, aggregated, and inserted into the appropriate database. The following services periodically scan the Long Term Canister for session data to aggregate and insert into the appropriate Tealeaf databases for reporting purposes:

- **Data Collector Service** - The Data Collector aggregates session data for use in Portal reports. This service can be monitored through logs. See “Data Collector Status” on page 94.
- **Visitor Database Extractor** - If you have licensed and installed IBM Tealeaf cxResults, the Visitor Database Extractor periodically scans the LTC for session data to insert into the Visitor database. See “Data Collector Status” on page 94.

Database-Sourced Portal Reports

The easiest mechanism for reviewing the status of Tealeaf databases is to check the Portal reports that use Tealeaf database data.

The following report types are useful for verifying that data is being properly inserted into the Tealeaf databases and reported through the Portal.

1. **Report Builder** - You can configure reports to monitor Tealeaf-defined events, total hits and total sessions, portal audit events, and visitor activities.
   - Any report created in the Tealeaf Report Builder can be inserted onto a dashboard. See "Tealeaf Report Builder" in the *IBM Tealeaf Reporting Guide*.
2. **Scorecards:** Scorecards provide summary information on key performance indicators (KPIs) and processes.

   Note: Access to Tealeaf scorecards requires IBM Tealeaf cxView, a separately licensable component of the IBM Tealeaf CX system. please contact your IBM Tealeaf representative.
   - You can use the available scorecards to identify if recent data is being collected to populate them.
   - You can drill down into scorecards to search for and review individual sessions to verify data against the configured scorecard. See "Using Scorecards" in the *IBM Tealeaf cxView User Manual*.
   - For more information on configuring scorecards, see "Configuring Scorecards" in the *IBM Tealeaf cxView User Manual*.
3. **Dashboards:** Dashboards enable the display of Tealeaf system data and data extracted from the Tealeaf databases in a widget-based display.

   Note: Access to Tealeaf dashboards requires IBM Tealeaf cxView, a separately licensable component of the IBM Tealeaf CX system. please contact your IBM Tealeaf representative.
   - Configured dashboards are available through the Dashboards menu in the Portal. See "Using Dashboards" in the *IBM Tealeaf cxView User Manual*.
   - You can use system dashboards to monitor activity and individual page information. For more information on system dashboards, see “System Dashboards” on page 84.
   - You can configure dashboards to display desired page and activity content. See "Configuring Dashboards" in the *IBM Tealeaf cxView User Manual*.
4. **Database Status Reports** - Through the System Status page, you can access two reports that provide detailed information on the status of individual databases and data groups.
   - See “System Status” on page 8.
   - See “System Status” on page 8.

Portal Database Queries

Through the Portal, you can monitor interactions between the Portal application and the Tealeaf databases.

- Some Portal queries are managed through an intermediate service, which is not available for testing through this page. See "Configuring the Tealeaf Data Service" in the *IBM Tealeaf CX Configuration Manual*.

Through the Portal Management page, you can access query information for the Portal application to each database and can test the AJAX toolkits in use by the Portal.

Note: The database queries information applies only to your current Tealeaf session. The SQL calls to each database can be useful for diagnosing database query issues. It is cleared whenever you log out of the Portal.

1. Login to the Portal as an administrator.
2. From the Portal menu, select **Tealeaf** > **Portal Management**.

3. Click the Logs link in the left navigation panel.
4. In the Logs pane, click the Portal History link.
5. From the View drop-down, select one of the following options:

Option Description

AJAX Call History
Information on AJAX calls made by the Portal application. You can test functionality of the AJAX toolkits used by the Portal through the Portal information report. See “Portal Application Information Report” on page 78.

System Database History
Review recent SQL calls made to the TL_SYSTEM database.

Reports Database History
Review recent SQL calls made to the TL_REPORTS database.

RSE Database History
Review recent SQL calls made to the TL_RSEXTRACTOR database.

Statistics Database History
Review recent SQL calls made to the TL_STATISTICS database.

- For more information on these databases, see "Overview of Tealeaf Databases" in the *IBM Tealeaf Databases Guide*.
- For more information on these reports, see “Portal Logs” on page 227.

Data Collector Status

Through the Portal, you can monitor the status of the Tealeaf Data Collector, which is responsible for collecting session data from the Long Term Canister and inserting it into the appropriate database.

You can review the current collection status through the Portal. **Steps:**

1. Login to the Portal as an administrator.
2. From the Portal menu, select **Tealeaf** > **System Status** > **Collection Status**.
3. The Collection Status report shows the most recent data collection runtime.
   - See “System Status” on page 8.

Through the Portal Management page, you can troubleshoot issues with Data Collection by reviewing the application logs and Windows Event logs.

- See “Troubleshooting Tealeaf CX Issues through the Portal” on page 97.
- The Data Collector logs are available for review and download through the Portal Management page. See “Portal Logs” on page 227.

Visitor DB Extractor Status

Embedded in the Tealeaf Data Service, the Visitor Database Extractor is used to collect sessions from the Long Term Canister and inserting them into the Visitor database.

Note: The Visitor database is a component of IBM Tealeaf cxResults. IBM Tealeaf cxResults is no longer available as a newly licensed product as of Release 8.7. Customers that licensed IBM Tealeaf cxResults in Release 8.6 and earlier may continue to use and receive support for the product in Release 8.7 and later. For more information, please contact Tealeaf Customer Support .

Through the Portal Management page, you can troubleshoot issues with the Visitor Database Extractor by reviewing the logs.

- See “Troubleshooting Tealeaf CX Issues through the Portal” on page 97.
- The following log filenames are applied to the daily log file for the Tealeaf Visitor Database Extractor: `TL_VDB_Extractor_YYYYMMDD.log`.

You can review logs from the Data Service through the Portal Management page. In the Logs panel, select the Data Service link.

- See “Portal Logs” on page 227.

Tealeaf Database Manager

The Tealeaf Database Manager is an external utility for installing, removing, or updating individual Tealeaf databases or all known databases in the system.

Through the Tealeaf Database Manager, you can review configuration information and metrics on each Tealeaf database.

- The Tealeaf Database Manager maintains its own set of logs and displays useful information when executed actions fail.
- See "Tealeaf Database Manager Reference" in the *IBM Tealeaf Databases Guide*.

Reporting on Integration

This section describes the integration of IBM Tealeaf.

cxConnect for Data Analysis

To monitor IBM Tealeaf cxConnect for Data Analysis, the following information resources are available.

- To open IBM Tealeaf cxConnect for Data Analysis, select **Tealeaf** > **IBM Tealeaf cxConnect for Data Analysis** from the Portal menu.

  Note: IBM Tealeaf cxConnect for Data Analysis is a separately licensable component of the IBM Tealeaf CX system. please contact your IBM Tealeaf representative.

1. **Scheduled Tasks:** In IBM Tealeaf cxConnect for Data Analysis Scheduled Tasks page, you can review the status of currently configured tasks. See "cxConnect Scheduling Tasks" in the *IBM Tealeaf cxConnect for Data Analysis Administration Manual*.
2. **Logs:** Every task listed in the Scheduled Task page generates a runtime log. To view this log, right-click the task and select **View Log**.
   - All IBM Tealeaf cxConnect for Data Analysis tasks are logged into the following file. This log file contains task information including date/time and the number of sessions each task extracted:

     `<Tealeaf_install_directory>\Logs\TLExtractJobHistory_YYYYMM.csv`
3. * Tealeaf Status report:* The Tealeaf Tealeaf Status report contains status information on IBM Tealeaf cxConnect for Data Analysis. See “Tealeaf Status Report” on page 89.

For more information on IBM Tealeaf cxConnect for Data Analysis, see "cxConnect for Data Analysis Administration Manual" in the *IBM Tealeaf cxConnect for Data Analysis Administration Manual*.

cxVerify

To monitor IBM Tealeaf cxVerify, the following information resources are available.

- To open IBM Tealeaf cxVerify, select **Tealeaf** > **IBM Tealeaf cxVerify** from the Portal menu.

  Note: IBM Tealeaf cxVerify is a separately licensable component of the IBM Tealeaf CX system. IBM Tealeaf cxVerify is no longer available as a newly licensed product as of Release 8.7. Customers that licensed IBM Tealeaf cxVerify in Release 8.6 and earlier may continue to use and receive support for the product in Release 8.7 and later. For more information, please contact Tealeaf Customer Support .

1. **Scheduled Tasks:** In IBM Tealeaf cxVerify Scheduled Tasks page, you can review the status of currently scheduled tasks. See "cxVerify Scheduling Tasks" in the *IBM Tealeaf cxVerify Administration Manual*.
2. **Logs:** Every task listed in the Scheduled Task page generates a runtime log. To view this log, right-click the task and select **View Log**.
   - All IBM Tealeaf cxVerify basic runtime information stats are logged into the following file. This log file contains task information including date/time and the number of sessions extracted by each task:

     `<Tealeaf_install_directory>\Logs\TLExtractJobHistory_YYYYMM.csv`
3. * Tealeaf Status report:* The Tealeaf Tealeaf Status report contains status information on IBM Tealeaf cxVerify. See “Tealeaf Status Report” on page 89.

For more information on IBM Tealeaf cxVerify, see "cxVerify Administration Manual" in the *IBM Tealeaf cxVerify Administration Manual*.

PCA Administration Reporting

Tealeaf administrators who are also monitoring the IBM Tealeaf CX Passive Capture Application and its host can use the links below to access monitoring tools and methods for the PCA.

Network Traffic

Many PCA-related issues can be identified by monitoring the network traffic into and out of the PCA. At the command line, you can use the Linux utility `tcpdump` to perform many of these monitoring tasks.

- See "Troubleshooting - Capture" in the *IBM Tealeaf Troubleshooting Guide*.
- Additional commands are available in the Web Console section. See "Passive Capture Configuration via Web Console" in the *IBM Tealeaf Passive Capture Application Manual*.

Reports generated by `tcpdump` are best reviewed through a viewer such as wireshark.

- For more information, see `http://www.wireshark.org`.

PCA Application

The IBM Tealeaf CX Passive Capture Application provides a variety of useful data on capture activities through the PCA Web Console.

- On the Summary page, you can review basic summary information about the capture process, as well as enable or disable the capture process. See "PCA Web Console - Summary Tab" in the *IBM Tealeaf Passive Capture Application Manual*.

- In the Statistics page, you can review configuration and resource usage information for the capture software processes and metrics on PCA performance. See "PCA Web Console - Statistics Tab" in the *IBM Tealeaf Passive Capture Application Manual*.

Additional PCA information is available through the Tealeaf Status report. See “Tealeaf Status Report” on page 89.

stats2csv

Statistics logs from the PCA can be converted to `.csv` or Excel files in `.xml` format using the `stats2csv` script in the PCA installation.

Monitoring Capture Status through TMS

Through TMS, you can monitor statistical information on the socket capture process for each pipeline in your environment.

- To access TMS, select **Tealeaf** > **TMS** from the Portal menu. See Chapter 4, “Tealeaf Management System,” on page 259.
- In the Pipeline Status tab, you can see the active connections, byte transfer rates, and page views as currently sampled in the pipeline data.
- You can also review how the data is passed through each session agent in the pipeline.
- See Chapter 6, “TMS Pipeline Status Tab,” on page 297.

Pipelines can be created and edited through TMS. See Chapter 8, “TMS Pipeline Editor,” on page 307.

Monitoring Capture Status through System Statistics

Optionally, the PCA can be configured to send statistics hits through the Windows pipeline for insertion into the Statistics database. These hits include useful statistical information, as displayed in the Statistics tab.

- See “System Statistics” on page 39.
- For more information on PCA stats, see "PCA Web Console - Statistics Tab" in the *IBM Tealeaf Passive Capture Application Manual*.

Troubleshooting Tealeaf CX Issues through the Portal

The sections below describe how to identify the activities of the main processing pipeline through the Tealeaf Portal:

- In some cases, additional troubleshooting may be required outside of the Portal.
- For a methodology for testing your Tealeaf solution from end-to-end, see "Testing Your Tealeaf Solution" in the *IBM Tealeaf CX Configuration Manual*.

Monitoring Active Status

In the Portal, the Active menu contains useful tools for assessing the current activities in Tealeaf. Review the menu items in the following order:

- **Active Status:** Displays status information on current activities on the web application.
  - Verify that the Generation time statistics seem valid.
  - Check the Error counts. Large error counts can be investigated by creating a report to track the application error. See "Tealeaf Report Builder" in the *IBM Tealeaf Reporting Guide*.

  - For more information on active status, see "Monitoring Active Status" in the *IBM Tealeaf cxImpact User Manual*.
- **Active Sessions:** Displays information on active sessions. If there are no active sessions, go to the monitored web application and begin browsing to create one.
  - There may be a delay of 90 seconds on average before your active session begins appearing in this report.
  - Be sure to click **Refresh** or select **Refresh Automatically**.
  - If you still see 0 active sessions after creating an active session, then see "No Current Activity" on page 99.
  - See "Monitoring Active Sessions" in the *IBM Tealeaf cxImpact User Manual*.
- **Active Events:** Displays information on the currently active events for the selected Canister.
  - This report may contain useful information on cumulative events over approximately the last 30 minutes. If you have created events to occur on every page or on every session, you can look for those events to gauge basic activities.
  - These canister events are accessed through Search Server and bypass the Tealeaf database. If there is no data in this report, then there may be a problem with Search Server. See "Troubleshooting - Searching" in the *IBM Tealeaf Troubleshooting Guide*.
  - If you have created events to monitor specific application errors or issues with the monitored web application, you may find the Watch feature useful.
  - To review active sessions containing the event, click the hyperlink number in the Live Count column.
  - To see all active events as a list, select `List` from the View drop-down.
  - See "Monitoring Active Events" in the *IBM Tealeaf cxImpact User Manual*.
- **Alert Monitor:** Displays status of the currently configured set of alerts.
  - The Alert Summary panel indicates Warning and Alert Threshold values as well as the Current value, which may be useful for identifying either significant problems in the system, depending on the alert, or a mis-configured alert.
    - When an alert is selected, search links in the displayed report provide access to active sessions or recently completed sessions that contain events that triggered the alert.
    - Recent alert history is available for download.
  - Since alerts are triggered based on events, you can create events to report useful information on the current state of the page or session.
  - To view a dashboard of the currently active alerts, click **View Gallery**. You can drill into these summary reports on each alert for additional detail.
  - See "Alert Monitor" in the *IBM Tealeaf cxImpact User Manual*.

IF:

- Above steps indicate that there is no current activity, then see "No Current Activity" on page 99.
- Some hits are reaching the Short Term Canister, then see "Troubleshooting Issues in the Canister" on page 101.

No Current Activity

If you have been unable to detect any current Tealeaf activities, including no active sessions after you have generated one, it indicates that sessions are not reaching the Short Term Canister.

The first test is to verify through the TMS Pipeline Status tab that the Windows pipeline is processing hits.

1. Login to the Portal as an administrator.
2. To access TMS, select **Tealeaf** > **TMS**.
3. Click the Pipeline Status tab.
4. From the Server drop-down, select the Processing Server to test.
5. Select the appropriate pipeline from the Pipelines textbox.
6. In the upper-right corner, for the selected pipeline on the selected server, you can review the data flow through the individual session agents in the pipeline.
   - Data should be flowing consistently through all listed session agents.
   - Some session agents, such as Data Drop, remove hits from the pipeline. Among the list of session agents, the numbers in the Hit Count column may decrease.
   - The Queued column indicates hits that have been queued for processing for the listed session agent. Typically, this situation indicates some mis-configuration of the session agent. In such cases, you should review the session agent's configuration.
     – See "Overview of the Capture Pipeline and Session Agents" in the *IBM Tealeaf CX Configuration Manual*.
7. Check the other pipelines on the selected server, if applicable.
8. Check the pipelines on other servers in the environment, if applicable.

IF

- Hits are getting bottlenecked with a specific session agent, then review the session agent's configuration. See "Overview of the Capture Pipeline and Session Agents" in the *IBM Tealeaf CX Configuration Manual*.
- No hits are passing through the pipeline, then “No Hits in the Pipeline.”
- Some hits are passing through the pipeline but not expected volumes, then “Some hits are arriving into the Canister” on page 103.

Troubleshooting Issues in the Pipeline

The Windows-based pipeline is composed of a sequence of individual agents that perform manipulations on hits that have been captured by the IBM Tealeaf CX Passive Capture Application and forwarded to the Tealeaf Transport Service.

- To analyze current processing activities of any configured pipeline, use the Pipeline Status tab. See Chapter 6, “TMS Pipeline Status Tab,” on page 297.
- Pipelines are created and edited through the Pipeline Editor. See Chapter 8, “TMS Pipeline Editor,” on page 307.

No Hits in the Pipeline

If there are no hits passing through the pipeline, then you must validate the upstream data source. Depending on your Tealeaf deployment, the upstream data source may be one of the following:

IF

- **IBM Tealeaf CX Passive Capture Application:** If the pipeline's source of data is the IBM Tealeaf CX Passive Capture Application, additional verification is required in the PCA Web Console, which is outside of the Portal. See “PCA.”
- **Health-Based Router:** In multi-canister environments, HBR can be used to manage load balancing. See “HBR” on page 101.
- **Other pipeline:** If the pipeline is a child pipeline, then you should review the configuration for passing hits to the child pipeline and the status of the parent pipeline. See “No Current Activity” on page 99.

Pipeline Data Sources

This section describes the pipeline data sources.

PCA

To verify that the PCA is successfully capturing TCP packets, evaluating the data, and forwarding hits to downstream consumers, you must verify that capture is properly configured and enabled in the PCA Web Console.

- **Console Tab:** Verify that capture is enabled and is capturing packets from the network.
  - See "PCA Web Console - Console Tab" in the *IBM Tealeaf Passive Capture Application Manual*.
- **Summary Tab:** Review the following:
  1. Machine Health section identifies each enabled PCA process and the current data volume and CPU usage.
  2. Peers section indicates the destinations to which the PCA is delivering hits.
  3. Current Per Second Stats section indicates the current processing load for individual processes in the PCA. The Reassd Out column indicates how much data is emerging from the PCA right now.
  - See "PCA Web Console - Summary Tab" in the *IBM Tealeaf Passive Capture Application Manual*.
- **Delivery Tab:** This tab can be used to troubleshoot delivery issues to the target recipients.
  - Target recipients should include the pipeline on the server on which you have noticed limit or no hit processing.
  - You can use the Ping link to test the connection between the PCA and the target.
  - Verify that the setting in the Security column matches the expected delivery protocol for the pipeline in question.
  - For issues in slow processing, the Speed test can verify the transfer rate across the network between the PCA and the target.
  - See "PCA Web Console - Delivery Tab" in the *IBM Tealeaf Passive Capture Application Manual*.
- **Interface Tab:** Use this tab to verify all instances of the IBM Tealeaf CX Passive Capture Application and to review the configurations for all network interface cards (NICs) that are listening to web traffic.
  - See "PCA Web Console - Interface Tab" in the *IBM Tealeaf Passive Capture Application Manual*.

For more information on testing your PCA configuration, see "Initial PCA Configuration" in the *IBM Tealeaf Passive Capture Application Manual*.

HBR

Through the Portal, you can review the current status of the HBR machine.

- Health-Based Routing enables load balancing between multiple Canisters in a Tealeaf environment. See "Health-Based Routing (HBR) Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

1. In the Portal menu, select **Tealeaf** > **System Status** > **Health Based Routing**.
2. In the General section, review the Health message.
   - If all Canisters are reporting good health, the message should read something like the following:

     `4 out of 4 Canisters Healthy and Available`
3. To identify if hits are being spooled, review the value for `Current Spooled Hits`.
   - This value must be assessed based on the current and expected traffic volume. You should divide this value by the number of hits per second. This value can be retrieved from the Pipeline Status tab in TMS. See Chapter 6, "TMS Pipeline Status Tab," on page 297.
4. Below the General report are two graphs indicating the percentage of healthy Canisters out of the total available and the percentage of traffic that is being properly passed to a Canister for evaluation.
5. At the bottom of the report, you can review the status information for individual Canisters configured to receive hits from HBR.
   a. For a health Canister, the Health State message should be the following:

      `Healthy`
   b. Other status messages typically include a `Reason`, which may be explored through other System Status reports, as described in the following sections. See "Some hits are arriving into the Canister" on page 103.
   c. If a Canister is overburdened and unable to keep up, you should verify that the Routing value does not indicate that the Canister is receiving too high of a percentage of traffic from HBR.
6. To refresh the report, click **Refresh** in the left navigation panel.

For more information on this report, see "System Status" on page 8.

IF:

- There is no data in the HBR report, then the HBR machine may be down.
  - In the Portal Management page, you should ping the machine hosting HBR. See "Managing Tealeaf Servers" on page 115.
- Some Canisters are reporting that they are down or in an unhealthy state, you should investigate each one using the Canister status report. See "Canister Health" on page 102.
- All Canisters are down, then HBR is unable to communicate with them.
  - In the Portal Management page, you can verify the configuration and availability of each Processing Server. See "Managing Tealeaf Servers" on page 115.
  - The problem may be an error in how the Canister configurations for HBR. See "Health-Based Routing (HBR) Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

Troubleshooting Issues in the Canister

This section describes troubleshooting issues in the canister.

Canister Health

The System Status report for Canisters provides diagnostic data for the Short Term and Long Term Canisters in the Tealeaf environment.

Steps: To access the Canister health report, please complete these steps.

1. In the Portal menu, select **Tealeaf** > **System Status** > **Canister**. The report for all Canisters is displayed.
2. Review the counts of Total Sessions in the STC and the LTC.
3. Review the number of Active Sessions. You should be able to find these sessions in the Active Sessions page. Select **Active** > **Sessions** from the Portal menu. See "Monitoring Active Sessions" in the *IBM Tealeaf cxImpact User Manual*.
4. The Un-indexed Sessions count can indicate a problem with indexing configuration. A high number of these sessions compared to the Total Sessions (LTC) figure should be investigated. See “Long Term Canister” on page 108.
5. A non-zero number of Non-Evaluated Hits indicates that one or more of the Canisters is spooling. You should review the details in individual Short Term Canisters. See “Short Term Canister.”
6. On the right side of the All Servers report, you can see the following statistics:
   a. Memory In Use indicates the percentage of all Short Term Canister memory allocated for session data that is currently in use for all servers.
      - If this figure is consistently over 80%, there are inadequate system resources for processing the data volume in the Short Term Canisters. You need to explore ways to free up system resources on your existing server or servers or consider adding resources. For more information, see your Windows Server documentation.
   b. CPU Utilization metrics should daily rates of CPU utilization and recent rates.
      - If this figure is consistently near 100%, there are inadequate system resources for processing the data volume in the Short Term Canisters. You need to explore ways to free up system resources on your existing server or servers or consider adding resources. For more information, see your Windows Server documentation.

For more information on this report, see “System Status” on page 8.

Short Term Canister

One component of a Tealeaf Canister, the Short Term Canister is an in-memory database used for storing active sessions.

To review information on an STC, in the Canister Status report select the individual Canister that you wish to review from the IBM Tealeaf CX drop-down. The following areas of the report can be reviewed to check STC health.

- Determine if hits are even arriving in the Canister. See “No hits arriving into Canister” on page 103. If some hits are arriving, continue with these steps
- **`Session Timeout`**. This value in seconds defines how long idle sessions are permitted to remain open in the Short Term Canister without updates or changes applied to them.
  – For more information on changing this setting, see "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

- **Sessions waiting for archive to disk**. This value indicates the number of sessions that have been marked for closing but have not been written to disk in the Long Term Canister. See "Sessions waiting for archive to disk."
  - A high number for this value typically indicates disk storage problems, I/O problems with the storage media, and processing resource issues.
    - Disk storage: Verify the amount of free disk space available under Disk Usage: `Server Free Disk Size`.
    - I/O problems: To verify I/O problems, use your external disk monitoring tools.
  - processing resource issues: It's possible to increase the number of session processors in the STC. However, adding these threads does cost available CPU and RAM resources. See "Hit Processors" on page 105.
- **Memory In Use**. If this figure is greater than 80% in most deployments, hits submitted to the STC for evaluation are spooled to local disk for later processing. See "Spooling" on page 105.
- **Non Evaluated Hits**. A non-zero figure indicates that the Short Term Canister is having difficulty processing events. See "Event processing" on page 107.

For more information on this report, see Single Server under "System Status" on page 8.

No hits arriving into Canister

To verify that there are no hits arriving in the Canister, do the following:

1. In the DecoupleEx status report, look at the following values under the Canister heading:
   - `Hits/Sec In Current`
   - `MB/Sec In Current`
2. If both of the above values are `0` and unchanging over a period of a few seconds, then there are no hits arriving in the Canister, and the system is spooling.
   - See "Spooling" on page 105.

Some hits are arriving into the Canister

The Canister Status report provides detailed information on the active and completed sessions stored in the Canister, as well as indexing operations.

- See "Canister Health" on page 102.

IF: If an incomplete set of hits is arriving into the Canister, then there may be a problem in how the Portal application is communicating with the Canister machine or the database.

- Data from the Canister may not be surfacing in the Portal reports.
- Data in the Canister may not be collected, aggregated, and inserted into the Tealeaf databases. See "Reviewing Data Collection Events and Logs" on page 111.

Then, there's something wrong with the Portal application talking to either the Canister machine or the database. *Not as frequent

Sessions waiting for archive to disk

If the Canister Status report indicates that there are a high number of `Sessions waiting for archive to disk`, then the Short Term Canister is unable to write to the Long Term Canister at a sufficient rate.

In most cases, this issue is caused by a basic I/O throughput problem. The incoming session count is higher than the outgoing session count. You can verify this issue through the DecoupleEx status report.

To fix this issue, you should identify if there is excess I/O capacity available using a tool such as SQLIO while Tealeaf is running.

- If there is excess capacity, you can experiment with increasing the number of Session Processors, whose primary function is to compress sessions for storage on disk. By increasing the processing units, more sessions can be written to disk at a faster rate. After raising this number, you may notice that the rate at which sessions are processed is increasing. For more information, see the `Session Processors` value in the "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.
- If there is no excess capacity, there is no software solution for this issue. Please review your storage configuration and capacity.

High ctree Memory

If the Canister message indicates `High ctree Memory`, the configured RAM aside for the STC is insufficient for current processing load. Do the following:

- Review your STC memory settings. See "STC memory settings."
- Identify if this issue has occurred a peak period or is a normal load
- Identify if there are any reasons why the load may have recently changed. For example, if recent changes to the web application have resulted in large average page sizes or more hits per session, STC processing capabilities may be affected.
- It may be that there are a high number of sessions in the STC waiting for writing to LTC. See "Sessions waiting for archive to disk" on page 103.

To resolve high ctree memory issues:

- Review your session timeout settings. The value `Session Timeout` in the Canister report indicates the current session timeout.
- Review your session size settings. Three parameters in the Services Controls tab define the limits in size of a session in terms of hits, time, or kilobytes of data.
  - For more information on changing the above settings, see "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.
- You should consider configuring Session Close events to terminate sessions earlier.
- Acquire more memory for the STC server.

STC memory settings

require-text

In the Canister report

- In most Tealeaf deployments, the STC begins spooling hits when `Memory In Use` reaches 80% of the total available memory, as defined in `Canister Memory Configuration`.
- **`Sessions waiting for evaluation (STC)`** indicates the number of Tealeaf sessions that have been spooled.
- Freeing up memory: You can review the following configuration settings to make more memory available in the STC for processing hits:

– **`Max Ctree Bytes`** In the Services Controls tab, this setting defines the amount of RAM memory available for the Short Term Canister. It should be evaluated based on other required resources on the Processing Server.
– **Session Size Limits** In the Services Controls tab, the following settings define the maximum permitted sizes of sessions, in terms of hits, data, and time. Reducing these size settings may increase the amount of free memory available for processing other hits.
  - `Max Hits Per Session`
  - `Max Size Per Session`
  - `Max Time Per Session`

  Note: The above settings may not be available in standard format. To change them, you may have to edit the Canister configuration in raw format. See “TMS WorldView Tab” on page 262.
– **Number of Hits Processors** In the Services Controls tab, you may choose to increase the number of Hit Processors. Some considerations apply. See “Hit Processors.”

Hit Processors

In most deployments, the Short Term Canister utilizes a single thread for processing hits. Depending on traffic volumes and available system resources, adding one or more Hit Processors can improve STC performance.

In the Canister report

- Compare the `Canister Memory Current Size` and `Canister Memory Max` values to the `Canister Memory Configuration` value, which indicates the maximum available memory for the STC. If the first two values are significantly less than the latter one, there should be available RAM to add another processing thread.
- Review the CPU Utilization statistics to identify if there is processing capability available for adding another thread.
- If resources are available, you can add a Hit Processor through the Services Controls tab in the Canister configuration in TMS.
- After adding the thread, you should continue to monitor the effects through this report.
  – See "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

Spooling

This section describes the spooling of session hits.

Detecting spooling

When the STC cannot process enough hits to match the incoming rate of hits, all new hits are written to a temporary file on the local disk until the memory usage has dropped to a sufficient level that both incoming and spooled hits can be safely processed.

- When the STC begins spooling, all hits from new sessions are spooled. Hits from sessions currently stored in memory continue to be processed through memory. For this reason, the percentage at which spooling begins should always be set below 100%.

In the DecoupleEx report

- Select each processing server.
- Check `Canister Status`. If it is not listed as `Healthy`, check the following:

– Canister Message options:
  1. High ctree Memory - The currently specified memory settings are inadequate for the load. See “High ctree Memory” on page 104.
  2. Non Evaluated Hits - If this figure is high, then the STC may be having difficulty processing events. See “Event processing” on page 107.

Spooling information is reported in the DecoupleEx System Status report. Key report values are described below.

- To access the DecoupleEx report, select **Tealeaf** > **System Status** > **DecoupleEx**.
- For more information on the related report, see “System Status” on page 8.

In the DecoupleEx report

- Select All Servers.
- Review the Canister Status value for each Canister.
  – A message of Spooling indicates that the Canister is spooling.
- The count of hits that have been queued is listed for each Canister:
  – Queued Hits in Memory - Queued hits in memory are waiting for evaluation. A high number may indicate CPU constraints.
  – Queued Hits On Disk - Queued hits stored on disk waiting for the opportunity to be evaluated.
  – Queued Hits Total - Sum of the two values above. This figure should be divided by the Canister Hit/Sec In value for the Canister to determine the number of seconds of queued data.
  – Hit/Sec In versus Hit/Sec Out - Difference between these two numbers indicates the rate of change in queuing.
    - If this value is > 0, then additional hits are being spooled.
    - If this value is < 0, then there is capacity to begin processing queued hits.

Spooling by Server

require-text

In the DecoupleEx report

- From the drop-down, select the same Processing Server where the spooling is occurring.
- Most of the information in this report is self-evident. Some key information:
  – DecoupleEx Status - When Extended Decoupler session agent is deployed and operating properly, this value is Memory.
  – Canister Status - When the Canister is operating properly, this value is Real-Time.
  – Canister Message - When the Canister is operating properly, this value is Canister is operational. Other Canister messages:
    - High ctree Memory - STC lacks sufficient memory to process the incoming rate of hits and is spooling. Find ways to free memory. See “High ctree Memory” on page 104.
    - Non Evaluated Hits - STC lacks sufficient resources to process events. See “Event processing” on page 107.

- `Total Dropped Hits` - This value should be `0` at all times. If it is not and spooling is in use, you should revisit the Extended Decoupler session agent configuration for this item to free up memory or processing and the hardware configuration of the server.
- `Current Spooled Hits` - This value indicates how many hits are currently stored on disk for the specified Canister.
- `Hits In/Out Ratio` - Rate of hits entering the STC divided by the rate of hits being written to the LTC
- `Time of First Hit in Spool Files` - Timestamp for when the current spooling of the Canister was initiated. Use this value to determine for how the Canister has been spooling.

Configuration options to limit spooling

Spooling is initiated and managed by adding the Extended Decoupler Session Agent to each pipeline. The following checklist identifies the settings that enable and configure spooling behavior:

- **Spool Triggers:** `Canister Max % Memory Used` - When this percentage is reached, the session agent begins spooling hits to disk for the Short Term Canister and does not stop spooling until `Canister Min % Memory Used` is reached. The following settings can also trigger spooling:
  - `Max Unevaluated Canister Hits` - When this number of hits is in the STC waiting for evaluation, spooling begins. Spooling does not stop until `Min Unevaluated Canister Hits` is reached in the STC.
  - `Max Sessions Waiting for LTC` - When this number of sessions is in the STC waiting for writing to the LTC, spooling begins. Spooling does not stop until `Min Sessions Waiting for LTC` is reached in the STC.
  - `Max Sessions Waiting for IDX` - When this number of sessions is waiting for indexing, spooling begins. Spooling does not stop until `Min Sessions Waiting for IDX` is reached.
- **Spool Drain:** The rate at which the spool is emptied is determined by the following configuration settings for the Extended Decoupler Session Agent.
  - `Control Outflow Rate` - When set to `ON`, the flow of hits from the queue back to the Canister is moderated based on the following settings:
    - `*Outflow Control Mode` - Defines the method of how to control outflow of hits from the spool. For details, see "Extended Decoupler Session Agent" in the *IBM Tealeaf CX Configuration Manual*.
    - `*Max Outflow Hits Per Second` - Depending on the Mode above, this setting specifies the maximum outflow in hits per second.
    - `*Max Outflow Bytes Per Second` - Depending on the Mode above, this setting specifies the maximum outflow in bytes per second.

For more information on the session agent, see "Extended Decoupler Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

- If you do not wish to enable disk queuing, you can use the lightweight predecessor to the above session agent. When memory limits are exceeded with the Decoupler session agent, the hits are simply dropped. See "Decouple Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

Event processing

When hits arrive in the Short Term Canister, events are evaluated, and any resulting actions are taken.

In the Canister report

- Review the number of `Non Evaluated Hits`.
- Compare the above value to the rate at which hits are entering the Canister. This information is in the DecoupleEx report.
- If the non-evaluated hits constitutes more than a few seconds of data entering the Canister, the problem is usually associated with event processing. In most cases, the source of the issue is that events have been added or changed in the system to cause the processing backlog.
- The solution is to increase the number of Hit Processor or Session Processor threads if resources are available. See “Hit Processors” on page 105.

Long Term Canister

One component of a Tealeaf Canister, the Long Term Canister is a database stored on the local disk of the Processing Server to retain completed sessions and their related search indexes for a configurable length of time.

To review information on an LTC, in the Canister Status report select the individual Canister that you wish to review from the IBM Tealeaf CX drop-down. The following areas of the report can be reviewed to check LTC health.

- `Un-indexed sessions` - This value indicates the number of sessions in the LTC waiting for indexing.
- `Total Sessions LTC` - Divide the above value by this one to identify the percentage of un-indexed sessions in the LTC.
  - If the above percentage is elevated for a significant period of time, there may be problems with session indexing. See “Session Indexing.”
- The `Memory In Use` and CPU Utilization statistics may provide insight into more general issues.

For more information on this report, see Single Server under “System Status” on page 8.

Session Indexing

When sessions are moved to the Long Term Canister, they are queued for indexing by the Session Indexer, which periodically polls the LTC for new sessions to index.

- For more information on indexing, see "Configuring CX Indexing" in the *IBM Tealeaf CX Configuration Manual*.

Index processing:

You can identify if there are potential problems with indexing.

1. Select then Canister Status report.
2. Select an individual canister.
3. Review the value for `Un-indexed sessions`.
4. The above value needs to be compared to the number of hits being submitted per second to the canister to effectively gauge if session indexing is a bottleneck.
   - The interval in which the Session Indexer checks for new sections is defined by default to be `10000` milliseconds (ten seconds). You can review this setting through TMS. See "Configuring CX Indexing" in the *IBM Tealeaf CX Configuration Manual*.

- If the number of un-indexed sessions is approaching the number of sessions to be inserted on average during the interval between index checks, Session Indexing may be having a difficult time keeping up with the traffic load.

To improve session performance, you can examine configuration settings in TMS.

Index storage:

In a typical deployment, generated indexes should be approximately the same size as the session data that they index.

- For more information on index storage metrics, see "Storage."

Storage

You can monitor storage usage on any and all Canisters in the system.

To review the storage report, select **Tealeaf** > **System Status** > **Storage**.

In the Storage Report

- **Review session storage:** For each server, the LSSN Space Used and LSSN Space Free columns indicate how space has been used and is available for storage of session data in the Canister. The percentage values may be most important.
- **Review index storage:** For each server, the Index Space Used and Index Space Free columns indicate how space has been used and is available for storage of indexes in the Canister. The percentage values may be most important.
  - On average, the ratio of session data storage to index data storage is 1:1.
- Index Style has a significant impact on system storage:
  - `FastRsp` - Standard indexing indexes only a selection of fields in the response.

For more information on the Storage report, see "System Status" on page 8.

Troubleshooting Issues in Reporting

After data has been stored in the Long Term Canister, the Tealeaf Data Collector periodically scans the LTC for new sessions to aggregate and insert into the Tealeaf databases, where the aggregated data is used to populate reports.

- By default, the Tealeaf Data Collector executes every 5 minutes.
- See "Data Aggregation and Retention" on page 241.

Due to connectivity issues, the Data Collector can fail to execute a scheduled run, or data load increases can cause the Data Collector to get behind in its collection. This section provides some troubleshooting steps for identifying issues in Tealeaf reporting and, by extension, the Tealeaf Data Collector.

Identifying if the relevant services are on

The Tealeaf Data Collector and the Data Service are responsible for retrieving data and inserting it into the appropriate Tealeaf database.

Please complete the following tests to verify that these services are operational.

- The Visitor Database Extractor polls the LTC for data to insert into the Visitor database. See "Troubleshooting Visitors Database Issues" on page 112.

1. In the Portal menu, select **Tealeaf** > **TMS**.
2. In TMS WorldView, select Servers from the View drop-down.
3. Select the server hosting the Data Collector.

4. The icon next to the Data Collector should be green.
   - If it is not, to start the Data Collector, select it and click **Start**. When it restarts, the icon becomes green again, and the Data Collector resumes polling the LTC at the appropriate interval.
5. If you have licensed IBM Tealeaf cxResults and have deployed the Visitor database, the Scheduling Service icon should be green.
   - If it is not, to start the Scheduling Service, select it and click **Start**. When it restarts, the icon becomes green again, and the Scheduling Service resumes polling the LTC at the appropriate interval. See "Troubleshooting Visitors Database Issues" on page 112.
6. If the appropriate services are working, then please continue with the next section.

Review database activity reports

Please complete the following steps to review the Tealeaf activity reports that pull data directly from the database.

1. In the Portal menu, select **Analyze** > **Report Builder**.
2. Change Focus to today's date, if it is not currently selected.
3. Through the Report Builder, you can open the activity reports. Click the Open icon. All activity reports begin with `Activity Report:` in the name.
   - `Traffic (MB)` - Indicates overall processing for the selected period.
   - `Round Trip` - Provides information on the round trip times for data for the selected period.
4. These reports are all powered by the Tealeaf database. If data is present in these reports, then the Data Collector is able to collect and store data.
   - For more information on activity reports, see "Tealeaf Report Builder" in the *IBM Tealeaf Reporting Guide*.

IF

- You do not see data in these reports for the currently selected date. Then, you should change the focus date to just the preceding date and re-review the reports. In this manner, you can step backward until you locate the date when the Data Collector stopped processing.
- You do see data in these reports, but the data seems to be low, continue with the next section.

Check Reports for Data

The above steps can only identify if the Data Collector is basically functioning. It may be possible that the Data Collector is behind in collecting data due to connectivity, traffic load, or other issues.

You can use event charts to monitor Data Collector performance. An event chart is generated by querying the database for the number of instances of an event for the specified time period. For example, you can generate an event chart for an event that occurs every session or on every page to provide indicate when the Data Collector is gathering data. Since event charts typically break down the reporting period into hourly periods, you can check to see when the most recent data is displayed in the event chart. Please complete the following steps.

1. In the Portal menu, select **Analyze** > **Report Builder**.
2. Change Focus to today's date, if it is not currently selected.

3. From the Event Selector, select an event that occurs every page or every session.
   - For example, if your web application uses a Session Start event or always begins with a Login page, a related event can be selected for this test.

     Note: Depending on traffic load and processing resources, creating an event that is triggered on every page may be problematic. If your system can handle the load, one example of an event triggered on every page is an event that is fired when the application message `HTTP - 200` is in the response. Since this message is part of every response, it is a valid event to test. For more information on creating events, see "RealiTea Viewer - Creating Events" in the *IBM Tealeaf RealiTea Viewer User Manual*.
4. After you have selected the event, click **Select**.
5. If the report does not automatically refresh, click **Refresh**
6. In the displayed report, compare the time for the most recently displayed bar in the bar graph to the local system time.
   - Any discrepancies between the hour in the bar chart and the current hour in your time zone indicates that there may be a general data flow problem in the main processing pipeline.
     – If you have already completed the preceding steps on this page and have not found the source of the issue, you may infer that the problem is with the Data Collector. Further analysis is required. See "Reviewing Data Collection Events and Logs."
     – You can compare the event count value in this report to the count among active events. See "Monitoring Active Events" in the *IBM Tealeaf cxImpact User Manual*.
   - If there are no differences, you might check the Count of events in the most recent hour in the displayed table to verify that the count is within reason.
   - To retrieve a list of sessions where the selected event was triggered within an hour, click the time indicator for the hour period.

Reviewing Data Collection Events and Logs

Through the Portal Management page, you can review the application events and associated logs for the Data Collector.

- The Portal Management tools are useful for diagnosing a variety of problems. This section provides diagnostics for the Data Collector, but similar steps can be applied to any Tealeaf service.

1. In the Portal menu, select **Tealeaf** > **Portal Management**.
2. Under Tealeaf Servers, click **Manage Servers**.
3. Select the server where the Data Collector is installed. By default, this server is the Report server.
4. When the server is selected, the server-specific toolbar is displayed at the top of the pane.
5. Please step through the following sections.

Reviewing timestamps

Through the Portal Management page, you can review timestamps for recent Data Collector and Result Set Extractor activities.

1. In the Portal Management page, select the Search Server whose data collection you wish to review.

2. With the Data Collector's server selected, click the Collector Dates tool in the toolbar ( )
3. A set of statistics is displayed. See “Managing Tealeaf Servers” on page 115.

Reviewing event logs

Through the Portal Management page, you can review Windows Event log entries for the selected server. These entries can provide detailed information on when and why the Data Collector failed.

1. With the Data Collector's server selected, click the Windows Event Log tool in the toolbar ( )
2. In the Event Log viewer, you can select the following values to begin your troubleshooting of the Data Collector:
   - Type: `Application`
   - Severity: `Error`
   - View: depends.
3. Review the list of Application errors to determine if any involve the Data Collector.
   - If you find specific errors, the log files may provide additional detail. See “Reviewing text logs.”
4. If you cannot find specific errors for the selected parameters, you can expand them to include non-Error events over a large time period, as necessary.
   - In rare cases, you may need to review `System` Type errors.

Reviewing text logs

For more detailed information on the activities of the Data Collector, you can access the logs through the Portal Management page.

- With the Data Collector's server selected, click the Tealeaf Logs tool in the toolbar ( )
- The logs for all Tealeaf services installed on that machine are available in the bottom panel.
- Use the drop-down to filter the display to show only the Data Collector logs.
- If you download a log, you should search the log for `ERROR` in caps.
  – In the Data Collector log, there may be fact names with `ERROR` in them. These may not be errors.

Troubleshooting Visitors Database Issues

If you have installed IBM Tealeaf cxResults, you can use the following sections to troubleshoot issues with the Visitors database.

Note: IBM Tealeaf cxResults is a separately licensable component of the IBM Tealeaf CX solution that creates a rich data repository on visitors' experience with your web application. IBM Tealeaf cxResults is no longer available as a newly licensed product as of Release 8.7. Customers that licensed IBM Tealeaf cxResults in Release 8.6 and earlier may continue to use and receive support for the product in Release 8.7 and later. For more information, please contact Tealeaf Customer Support .

Search

The Visitors database relies on the Visitor Database Extractor service to gather information for staging and insertion into the Visitor database.

- The Visitor Database Extractor is a standalone executable that is managed through the Tealeaf Scheduling Service. See "Configuring the Scheduling Service" in the *IBM Tealeaf CX Configuration Manual*.

Through the Visitor search page, you can see how current the data is in the Visitor database.

1. From the Portal menu, select **Search** > **Search for Visitors**.
2. In the Visitor Search page, compare the dates in the left navigation panel for session data and visitor data.
   - Depending on when the Visitor database was installed, the start dates may not match well.
   - The end dates should be within two days of each other. The Available Dates for Visitors should be no more than two days behind those for Sessions.

IF

- If the date range is more than two days:
  - Verify that the Scheduling Service is running through TMS. See “TMS WorldView Tab” on page 262.
  - The Visitor Database Extractor service may be having troubles. Review the logs for the service through the Portal Management page. See “Visitor Database Extractor log.”

Activity Reports

You can dig into recent visitor data through the activity reports.

1. In the Portal menu, select **Analyze** > **Report Builder**.
2. Change Focus to today's date, if it is not currently selected.
3. Click the Open tool in the toolbar. Click the Activity Reports group. Select the report to review.
   - You can use the DecoupleEx System Status report to compare these visitor counts to the number of one-hit sessions. If both values are high, there is a problem in how a session or visitor is detected.

Visitor Database Extractor log

Through the Portal Management page, you can review the individual logs of the Visitor Database Extractor

1. In the Portal menu, select **Tealeaf** > **Portal Management**.
2. Under Tealeaf Servers, click **Manage Servers**.
3. Select the server where the Report Server.
4. Click the View Logs icon in the toolbar ( ).
5. In the bottom panel, filter the display by `VDB Extractor`.
6. The displayed set of logs are for the most recent runs of the Visitor Database Extractor.
   - The timestamps indicate the date that the run was begun in GMT time.
7. If you download a log, you should search the log for `ERROR` in caps in the `TL_VDB_Extractor_YYYYMMDD.log` file.

About Page in the Portal

In the Portal, the About page provides significant information and a few tests of the Portal application.

- See “Portal Application Information Report” on page 78.

Next Steps

For more information on debugging specific issues with a component of the Tealeaf system, see "Tealeaf Troubleshooting Guide" in the *IBM Tealeaf Troubleshooting Guide*.

Chapter 3. Portal Management

This section contains administration topics for managing the Tealeaf Portal and features available through it.

In the Portal, you can manage user accounts, user groups, and settings for installed Tealeaf components. Through menu-driven controls, you can also configure reports, search templates, and browser-based replay, a web-based application for replaying visitor sessions. You can also create announcements that are displayed in the Portal.

To access:

1. Login to the Portal as an administrator.
2. Select **Tealeaf** > **Portal Management**.
3. The Portal Management page is displayed.

Note: After saving changes in the Portal Management page, you may need to logout of the Portal and login again to see the effects of those changes.

Use the following table to see more information about the Tealeaf Portal:

Table 4. More information about the Tealeaf Portal

For more information about...	See...
Tealeaf Portal	"Tealeaf Portal" in *IBM Tealeaf CX Basics*
Portal features	"Portal Features" in *IBM Tealeaf CX GlossaryIBM Tealeaf CX Glossary*
Portal search	"Searching in the Portal" in *IBM Tealeaf cxImpact User Manual*
Logging into the Tealeaf Portal, using Tealeaf Portal, Portal navigation	"Logging in to the Portal", "Using the Tealeaf Portal", and "Portal Navigation" in *IBM Tealeaf cxImpact User Manual*
Configuring Tealeaf Portal	"Configuring the Tealeaf Portal" in *IBM Tealeaf CX Database Guide*
Configuring Portal authentication	"Portal Authentication" in *IBM Tealeaf CX Configuration Manual*
Configuring Portal web application, configuring Portal for a locale	"Configuring the Tealeaf Portal" n *IBM Tealeaf CX Installation Manual*
Resolving any problems with the Portal	*IBM Tealeaf CX Troubleshooting Guide*

Managing Tealeaf Servers

In the Tealeaf Servers area, you can review all available Tealeaf servers, which may include IBM Tealeaf CX Transport Servers, IBM Tealeaf CX datastores and IBM Tealeaf CX indexing machines for search, event detection, and session replay.

The server definitions configured in the Portal Management page are used by the Portal for communicating with the servers and services in your Tealeaf system.

- To manage your Tealeaf servers, select **Tealeaf** > **Portal Management** > **Tealeaf Servers**.

© Copyright IBM Corp. 1999, 2014

- To display inactive canister servers, click **Display Inactive**. See “Deactivating a server” on page 125.

To view quick definitions of all Tealeaf servers, see "Tealeaf Servers" in *IBM Tealeaf CX Glossary*.

Managing Your Servers

This section provides information about managing your IBM Tealeaf servers.

To view quick definitions of all Tealeaf servers, see "Tealeaf Servers" in *IBM Tealeaf CX Glossary*.

Best Practices

The following sections provide recommended practices for how to configure Tealeaf servers through the Portal. Before you begin adding or editing servers, please review these sections.

Specifying Server Addresses

Depending on the networking environment in which the Tealeaf servers are deployed, you may be required to specify server addresses as IP addresses, domain names, or fully qualified domain names.

Any specified name must be resolvable by the Portal, which is hosted on the Portal Server, so that it can interact with the Tealeaf software hosted on the server.

Note: For additional requirements, please contact your network infrastructure administrator.

When specifying the address of a server, the following lists the order of preference for Tealeaf in how to specify it:

- Server names may be up to 256 characters in length.

1. **IP Address:** A direct IP address in the form of `255.255.255.255` for the machine hosting the Tealeaf server is the best solution.
2. **hostname:** Hostnames such as `myhost` or `server1` can be used, although there may be problems associated with DNS validation of the hostname.
   - In some environments, you may need to specify the hostname as a fully qualified domain name.
3. **`localhost`:** The `localhost` reference should be avoided, as it can be misinterpreted.

   Note: Avoid using `localhost` to specify a server hosted on the Portal Server. Exceptions:
   - The SQL Server can be referred to as `localhost`, if it is hosted on the same machine as the Portal.

Constraints on Hosting Tealeaf Servers

In general, you can specify a Tealeaf server to be co-hosted with any other Tealeaf server with the following constraints:

- For Canister servers, the entered host name cannot resolve to the host name, IP address, or full qualified domain name of another active or inactive Canister Server.

Note: When creating server entries, `tlserver.com` and `tlserver.domain.com` resolve to the same server and are not permitted. If you end up with multiple servers sharing a name as identified through hostname, fully qualified domain name, or IP address, you must deactivate one of them. To do so, edit the server to deactivate, deselect the Active checkbox and change the hostname of the server in a single edit. For example, `tlserver.domain.com` could be changed to `tlserverOLD.domain.com` in the inactive server. These multiple changes must be performed at the same time.

Adding a New Server

When you install the IBM Tealeaf CX system, all servers required for the Portal are automatically created for you. In most cases, you should not need to define new Tealeaf servers. As needed, however, you may add new Tealeaf servers.

- To add a new Tealeaf server, click **New**.

In some cases, new instances of required servers cannot be created, since the system needs one and only one server instance. These items do not appear in the New drop-down menu.

Note: You may not create duplicate instances of the Canister or Canister Search Servers on the same host, even if the existing instance is inactive.

- Some server definitions, such as the Capture Application Server, are not required for the Portal to function properly. However, by adding these servers to the Portal Management page, you may enable a set of commands through the displayed toolbar. See "Server Toolbar" on page 119.
- The Tealeaf Status Report may access servers through the Portal Management page definitions. See "Tealeaf Status Report" on page 231.

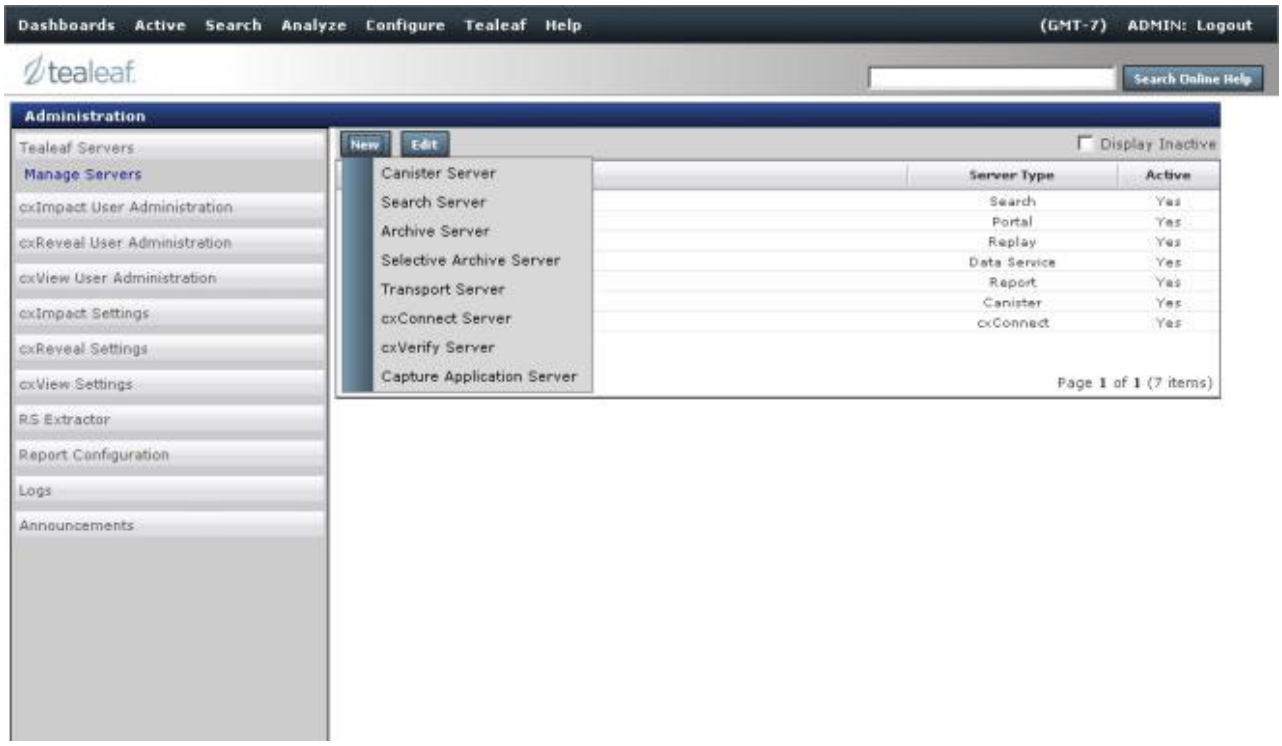

Figure 52. Adding a New Tealeaf Server

Server Description

Portal Server
The Portal Server entry specifies a reference to the server hosting the Portal application. This entry is created for you and does not normally need modification. See “Portal Server” on page 125.

Canister Server
A Canister Server is used for storing captured hits for active and completed sessions. You should consider adding a new Canister Server if captured hits are being spooled. See “Canister Server” on page 131.

Canister Search Server
Add a Tealeaf Canister Search Server, which provides session search capabilities to Tealeaf components and the IBM Tealeaf CX RealiTea Viewer (RTV) standalone application. See “Canister Search Server” on page 133.

cxReveal Database Search Server
If IBM Tealeaf cxReveal database search has been installed and configured, you must create an instance of the server hosting the IBM Tealeaf cxReveal database so that the users can perform searches of the database through the Portal. See “cxReveal Database Search Server” on page 125.

Replay Server
The Replay Server prepares and delivers sessions for replay through Browser Based Replay to requesting Tealeaf users. Your Tealeaf environment may contain multiple Replay Servers. See “Replay Server” on page 127.

Data Service Server
The Portal requires a reference to the server hosting the Tealeaf Data Service, which manages interactions with the Tealeaf datastores. See “Data Service Server” on page 130.

Transport Server
The Transport Server manages the Health-Based Router servers in the system. See “Transport Server” on page 134.

cxConnect Server
The IBM Tealeaf cxConnect for Data Analysis Server manages the assembly and delivery of session data to third-party applications. See “cxConnect Server” on page 136.

cxVerify Server
The IBM Tealeaf cxVerify Server is used to deliver session data for archiving in your enterprise data warehouse. See “cxVerify Server” on page 137.

Capture Application Server
The Capture Application Server is used to capture requests and responses for your web application. For the Portal, these servers provide statistical information on capture and pipeline operations. See “Capture Application Server” on page 137.

TLI Server
If you have deployed a TLI server to capture and manage static content for replay, you must create a reference to the server for the Portal application. See “TLI Server” on page 134.

Only one instance of each of the following servers is required for the Tealeaf solution. Depending on the Tealeaf products you have licensed, an instance is created during installation.

- These servers are not available in the New menu but can be modified if needed.

Server Description

Report Search Server
The Report Search Server is used for managing searches throughout the Tealeaf system, except for searches of the Visitor database. See "Report Search Server" on page 131.

Visitor Report Server
Identifies the location of the Query Server service that communicates with the IBM Tealeaf cxResults database. See "Visitor Report Server" on page 131.

Editing a Server's Properties

To edit a server's properties, select the server and click **Edit**.

The configuration for that server is displayed beneath the grid. The properties vary depending on the type of server. For more information on the properties for specific servers, use the links below.

- "Canister Server" on page 131
- "Transport Server" on page 134
- "cxConnect Server" on page 136
- "cxVerify Server" on page 137
- "Capture Application Server" on page 137

Server Toolbar

Depending on the type of server you have selected, a toolbar may appear with some or all of the following commands in it. The toolbar above the table contains buttons for all actions that can be performed on a server.

The example below appears when you select a Canister Search Server:

Figure 53. Canister Search Server Toolbar

For the currently selected server, the following toolbar commands may apply. They are listed below in left-to-right order as they may appear in the toolbar:

Table 5. Server Toolbar.

Icon	Command	Description
	Ping Server	Ping the server to test connectivity. See "Ping Server" on page 120.
	Collector Dates	Displays the dates when data was collected from the server. See "Collector Dates" on page 121.

Table 5. Server Toolbar (continued).

Icon	Command	Description
	Collector Statistics	This command displays the statistics of its most recent Data Collector run. See “Collector Statistics” on page 123.
	Collector Index	Contains information on the most recent indexing operations performed on the server. See “Collector Index” on page 123.
	Long Term Archive	Identifies the most recent operations to move session data to the long term archive. See “Long Term Archive” on page 123.
	Search Indexes	Indicates the most recent indexing operations for searching. See “Search Indexes” on page 124.
	Backup Info	Lists the most recent backup activities for the server. See “Backup Info” on page 124.
	Command Line	Opens a command line interface with the server. See “Command Line” on page 124.
	Registry and Services	Displays Windows registry and Services information for the server. See “Registry and services” on page 124.
	Windows Event Log	Displays Application or System event log messages in the Portal. See “Windows Event Log” on page 124.
	Tealeaf Logs	Displays all Tealeaf logs. See “Tealeaf Logs” on page 125.

Ping Server

Tests whether the Report Search Server is able to access the server hosting the IBM Tealeaf CX datastore and indexer as configured in the server definition. The ping returns version information for the database and Report Search Server.

Search Server Version
Reports the current version of the Search Server for queries

Search Server Alias
An alias to be used when an alternate hostname should be used for DNS resolution

Search Server NT Authorization
Tells whether NT Authorization is turned on the Indexer machine.

Search Server Auth Groups
Contains the name of the group to which a user must belong for normal Search Server commands.

Search Server Auth Admin Groups
Contains the name of the group to which a user must belong to perform administrative Search Server commands.

Search Server Auth Master
Authentication master server for the Search Server

Search Server Canister
Search Server's canister

Search Server Portal Server
The server hosting the Portal to which this Search Server connects

Search Server Alert Edit Enable
Ability to edit Alerts, either On or Off

Search Server Archive Server
The server to which the Archived sessions for this Search Server are transferred

Collector Dates

Displays the timestamps and the UTC ticks for recent collection history. This information is helpful to locate problems on the IBM Tealeaf CX server.

Entry Description

Data Collector Service Start
Most recent start or restart of the Data Collector service

Last Data Collection Run
Most recent run executed by the Data Collector

Last Interval for Scorecards (SID=5/Day)
Most recent day for which scorecard data has been calculated

Last Interval for Scorecards (SID=4/Week)
Most recent week for which scorecard data has been calculated. Week is reported as the first day of the week.

Last Interval for Scorecards (SID=3/Month)
Most recent month for which scorecard data has been calculated. Month is reported as the first day of the month.

Last Update for Scorecards
Most recent update for scorecards in the Portal

Last Reporting Data Trim
Most recent trimming of the report data stored in the database

Most recent hourly aggregation ? last hour
The latest hour aggregated in the most recent aggregation run of hourly reporting data

Most recent daily aggregation ? first day
The earliest day aggregated in the most recent aggregation run of daily reporting data

Most recent daily aggregation ? last day
The latest day aggregated in the most recent aggregation run of daily reporting data

Most recent hourly aggregation ? first hour
The earliest hour aggregated in the most recent aggregation run of hourly reporting data

Last Hourly Top Mover interval
Most recent hour of hourly Top Mover data

Last Daily Top Mover interval
Most recent day of daily Top Mover data

Last Hourly Top Mover calculation time
Most recent time that hourly Top Movers were calculated

Last Daily Top Mover calculation time
Most recent time that daily Top Movers were calculated

Last Calendar Table Load
Most recent time when the calendar table was loaded. This run occurs once per day.

Last Canister Data Trim
Time of most recent data trimming operation on canister data tables in the reporting database

Last Time of Daily Aggregation
Most recent time that daily reporting data was aggregated

Last Time of Dimension Log Aggregation
Most recent time of dimension log aggregation. These log runs are used to populate potential values for dimensions.

Last Dimension Log Interval
The latest hour aggregated in the most recent dimension log aggregation run. These log runs are used to populate potential values for dimensions.

Last Time of Hourly Aggregation
Last time that hourly reporting data was aggregated

Most recent hourly performance aggregation ? first hour
The earliest hour aggregated in the most recent aggregation run of hourly performance data

Most recent hourly performance aggregation ? last hour
The latest hour aggregated in the most recent aggregation run of hourly performance data

Most recent daily performance aggregation ? first day
The earliest day aggregated in the most recent aggregation run of daily performance data

Most recent daily performance aggregation ? last day
The latest day aggregated in the most recent aggregation run of daily performance data

Last Time of Daily Performance Aggregation
Last time that daily performance data was aggregated

Last Time of Hourly Performance Aggregation
Last time that hourly performance data was aggregated

Last RSE Aggregation Interval
Most recent hour interval of RSE event statistics aggregated

Last Time of RSE Aggregation
Last time of RSE event statistics aggregation

Minimum hour aggregated since last daily aggregation
Minimum hour interval of event data aggregated since the last successful aggregation of daily event data. Used for re-aggregation if aged data is captured.

Minimum hour aggregated since last daily performance aggregation
Minimum hour interval of event data aggregated since the last successful aggregation of daily performance data. Used for re-aggregation if aged data is captured.

Minimum hour aggregated since last hourly Top Mover calculations
Minimum interval aggregated since the last successful calculation of hourly Top Mover data. Used for re-calculating Top Movers if aged data is captured.

Minimum hour aggregated since last daily Top Mover calculations
Minimum interval aggregated since the last successful calculation of daily Top Mover data. Used for re-calculating Top Movers if aged data is captured.

Last Time of Dimension Value Tracking
Timestamp of when the dimension values were checked for most recent occurrence, which is used for dimension value trimming

Last Time of Dimension Trim
Timestamp of when dimensions were checked to see if values needed to be trimmed

- See "Data Management for Dimensions" in the *IBM Tealeaf Event Manager Manual*.

Collector Statistics

Shows collection times, number of read records, and length of time reading from the IBM Tealeaf CX Datastore:

Field Value

Collector
Collection process

Records
Number of records read.

Elapsed Time(s)
Elapsed time for last successful collection

Last Collection
Timestamp of last successful collection

Collector Index

Field Value

Index Category
Provides a list of table indexes from which the collector has gathered data.

Index Value
Represents the last record in the table, which provides a way for Tealeaf Support to know the location from which the collector last copied data.

Long Term Archive

Field Value

Name Name of the session data file.

Sessions
Number of sessions in the session data file.

First Use
Timestamp of first session in the session data file.

Last Use
Timestamp of the last session in the session data file.

Last Update
Timestamp of the last time the session data file was updated.

Offline Date
Expiration date of session data file.

Search Indexes

Reads the IXLIB.TLL file located on the indexing host machine and displays all indexes included in the file. The file is a master index that keeps track of all the indexes for sessions in the archive.

Backup Info

Checks the TBAK.DAT table in the archive and retrieves backup information, including the following:

- Date and time of last backup.
- Type of backup performed: full, partial, a specified date range.
- Location of the backup directory.

Command Line

The command line provides an alternative to browser-based interaction with Tealeaf servers. The textbox auto-completes commands.

- To run the specified command against the selected server, click **Execute**. The command's output is displayed in the box above the command line textbox.
- To see a list of all valid commands, click **View All Commands**.

Registry and services

You can view the current registry settings and services installed on the selected server. Select one of the following options of the following options from the View drop-down. Then, click **Refresh**.

Options
Value

Registry Settings
The Tealeaf registry settings on the server.

Services
The installed Tealeaf services. An icon shows the running state of each service.

Windows Event Log

Displays recent entries in the Windows Event Log on the selected server. The following settings are available.

- To apply your changes, click **Refresh**.

Setting Value

Type Displays either Application or System events.

Severity
Filters events by severity.

View The period over which to display events.

Tealeaf Logs

Displays a list of the log files generated by Tealeaf on the selected server.

- Each entry is a link to download the associated file.
- You can use the Filter By option to filter the display by log type.

Deactivating a server

If needed, you can remove a server from use by the Tealeaf system. The server information is not deleted; it is simply removed from usage. If needed, you can reactivate the server to resume operations.

- Servers are not deleted, as reports that reference the deleted server would no longer be functional.

  Note: Before you deactivate a server, you should verify that Tealeaf operations can continue without the server. Do not deactivate necessary servers that are the single active instance. For example, if you deactivate a Canister Server, Tealeaf users are unable to find any sessions stored on that server.

  To deactivate a server:

1. Select it, and click **Edit**.
2. In the display pane, deselect the Active checkbox.
3. Click **Save**.
4. The server is deactivated and is no longer used in Tealeaf operations.

Portal Server

The Portal Server manages the Tealeaf Portal web application, through which Tealeaf users access Tealeaf functionality.

- See "Using the Tealeaf Portal" in the *IBM Tealeaf cxImpact User Manual*.

Property
Value

`Use HTTPS`
If the Portal (IIS) is secured with HTTPS, this should be set to ensure URLs pointing back to the Portal are generated properly.

Note: If this value is selected, please verify that the Port number is configured appropriately.

`Display Name`
The name you want displayed under the Server heading on the IBM Tealeaf CX Server Configuration page.

`Host Name`
The DNS name or IP address of the Portal Server.

- The value for the Portal Server in this tab should match the value for the Portal Server in Search Server configuration. See "Configuring the Search Server" in the *IBM Tealeaf CX Configuration Manual*.

`Port` The port number to which the Portal listens.

- For non-HTTPS traffic, the default value is `80`.

cxReveal Database Search Server

The IBM Tealeaf cxReveal Database Search Server provides to the Portal the ability to search the IBM Tealeaf cxReveal database on the specified server.

Note: The cxReveal database is a component of cxReveal, a separately licensable component of the Tealeaf CX platform. please contact your IBM Tealeaf representative.

Note: The IBM Tealeaf cxReveal database and the ability to search it require separate installation and configuration.

- See "cxReveal Installation" in the *IBM Tealeaf cxReveal Administration Manual*.
- See "Configuring Session Attribute Search" in the *IBM Tealeaf cxReveal Administration Manual*.

During IBM Tealeaf cxReveal installation, a new Search Server entry is created for you in the Portal Management page. Additional configuration is not likely to be required.

- You can execute commands and download logs from the server through the Portal Management page. See “Server Toolbar” on page 119.
- The database that is searched may or may not be co-located on the same machine that hosts the IBM Tealeaf cxReveal server. For each server specified in the Portal Management page, you may specify a server hosting the IBM Tealeaf cxReveal database instance to query.

Note: To uninstall a IBM Tealeaf cxReveal Server that is hosted on a machine that is remote from the Portal Server, you must also remove the server reference from the Portal Management page after removing the software.

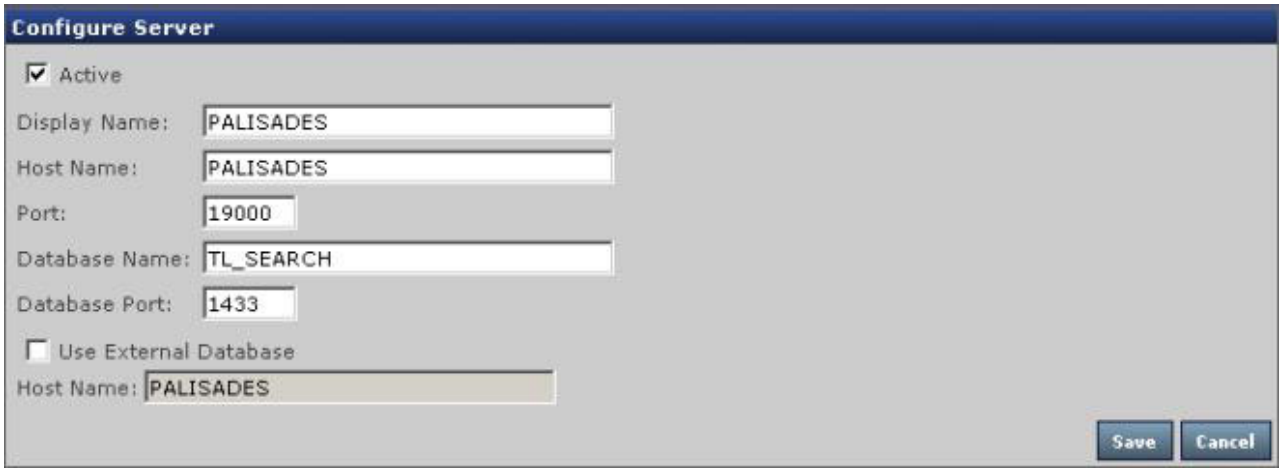

Figure 54. Configuring IBM Tealeaf cxReveal Search Server

Property
Value

Active
When set to Yes, the IBM Tealeaf cxReveal Search Server instance is active, and users may search it for sessions through the Portal.

Note: If you disable a IBM Tealeaf cxReveal Search Server and the database referenced in it is not referenced by another instance of the IBM Tealeaf cxReveal Search Server, then the database cannot be searched, and the sessions in it are unavailable to users.

Host Name
The DNS name or IP address of the IBM Tealeaf cxReveal Search Server. This value must be able to populated and resolvable by the server that hosts the Portal.

Note: An non-blank value is required. If the server is hosted on the same machine as the Portal application, you may enter `localhost`.

Port The port on which the IBM Tealeaf cxReveal Search Server instance listens to the Portal. Default is `19000`.

Database Name
The name of the IBM Tealeaf cxReveal database on the server.

- The default value is `TL_SEARCH`. If you are unsure of the name of the database, you can review it through the Tealeaf Database Manager. See "Using Tealeaf Database Manager" in the *IBM Tealeaf Databases Guide*.

Database Port
The port to use to communicate with the IBM Tealeaf cxReveal database.

- The default value is `1433`. If you are unsure of the name of the database, you can review it through the Tealeaf Database Manager. See "Using Tealeaf Database Manager" in the *IBM Tealeaf Databases Guide*.

Use External Database
When selected, queries to the specified IBM Tealeaf cxReveal Search Server are applied to a database instance hosted on the server on the specified host.

Note: In a typical deployment, the IBM Tealeaf cxReveal database is co-located on the server hosting the IBM Tealeaf cxReveal Search Server.

Host Name
When the Use External Database checkbox is selected, this value specifies the host server for the IBM Tealeaf cxReveal database.

Replay Server

The Replay Server is the core of Browser Based Replay, which allows Tealeaf users to replay captured sessions through their browser.

Note: If your web application is not compatible with the Mozilla family of browsers, Browser-Based Replay is unlikely to work. In Release 8.0 and later, the Replay Server utilizes the Gecko browser control to render session content on the server for delivery to the browsers of Tealeaf users. Replay of such sessions may still be supported through the IBM Tealeaf CX RealiTea Viewer, a desktop client that must be separately installed on each Tealeaf user's desktop. See "RealiTea Viewer (RTV) User Manual" in the *IBM Tealeaf RealiTea Viewer User Manual*.

In smaller Tealeaf user communities, only one Replay Server is required. However, if you are noticing performance impacts in Browser Based Replay, you may need to create additional Replay Servers in your Tealeaf environment to balance the load of delivering simultaneous replays to BBR users.

- See "Configuring the Replay Server" in the *IBM Tealeaf CX Configuration Manual*.

The following commands are available for the Replay Server.

Table 6. Replay Server

Icon	Command	Description
	Ping Server	Ping the server to test connectivity. See “Ping Server” on page 120.
	Server Stats	Collect and display statistics from the Replay Server. **Note:** After stats are displayed, click the Calendar tool to display stats from other dates. The other links are provided from Search Server and do not affect the displayed Replay Server stats.
	Renderer Control Stats	Collect and display statistics from the active renderers for the server.

To configure replay rules for the replay server, click **Rules** in the toolbar.

Note: You can configure replay rules only on the master Replay Server, which provides rules to all slave servers.

- See "BBR Replay Rules" in the *IBM Tealeaf cxImpact User Manual*.

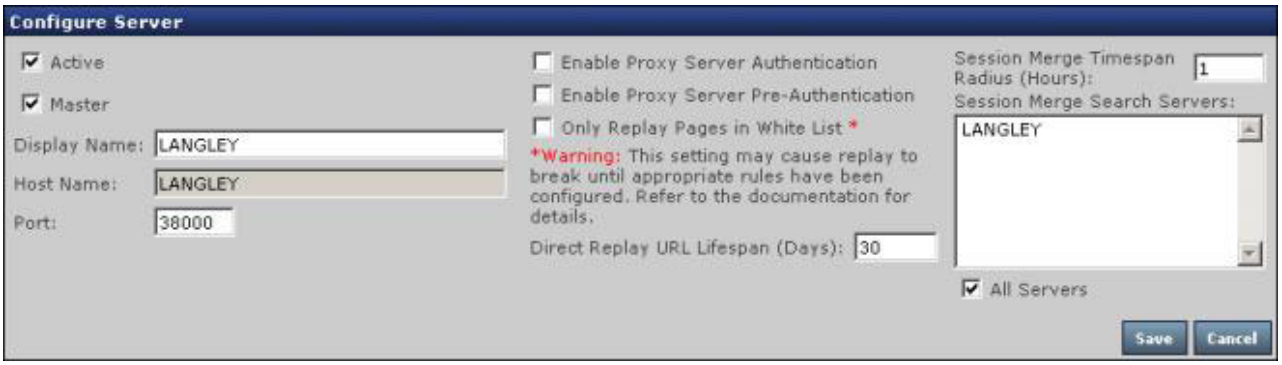
Configure Server
Active
Master
Display Name: LANGLEY
Host Name: LANGLEY
Port: 38000
Enable Proxy Server Authentication
Enable Proxy Server Pre-Authentication
Only Replay Pages in White List *
*Warning: This setting may cause replay to break until appropriate rules have been configured. Refer to the documentation for details.
Direct Replay URL Lifespan (Days): 30
Session Merge Timespan Radius (Hours): 1
Session Merge Search Servers:
LANGLEY
All Servers
Save
Cancel

Figure 55. Replay Server

The Replay Server has the following properties:

Property

Value

Active

When set to Yes, the Replay Server instance is active and able to serve sessions to BBR users.

Note: A master Replay Server is required. You cannot disable the only Replay Server in your environment. If you have multiple Replay Servers and disable the master, the Portal automatically designates another Replay Server as the master.

Master

When set to Yes, the Replay Server instance is designated as the Master Replay Server.

- One and only one Replay Server must be designated the Master server.

- See "Configuring the Replay Server" in the *IBM Tealeaf CX Configuration Manual*.

Display Name
The label used to identify the server.

Host Name
The DNS name or IP address of the Report Search Server. This value must be able to be resolved by the server that hosts the Portal. A valid non-blank value is required.

Port The port on which the Replay Server instance listens to the Portal. Default is 38000.

Enable Proxy Server Authentication
If the BBR user is behind a proxy server that requires authentication to connect to the outside, you can supply a username and password that BBR uses automatically when contacting the proxy. For example, "Proxy Server Authentication for Replay Server" on page 130.

Enable Proxy Server Pre-Authentication
The above setting uses a static set of credentials. With this setting, BBR makes a request through the proxy via the supplied URL before the session is loaded. With this authentication method, users only provide their credentials once, which the browser remembers, as opposed to multiple times when the session loads due to parallel HTTP requests. When enabled, enter the URL to the image database in the provided textbox.

Only Replay Pages in White List
When enabled, this setting configures Replay Server to replay only the URLs that have been added to the white list for the specified domain.

Note: When enabled, replay may not work until rules have been configured to specify the appropriate white list of URLs. See "BBR Replay Rules" in the *IBM Tealeaf cxImpact User Manual*.

Direct Replay URL Lifespan (Days)
This setting controls how long links into BBR are active after they are generated. When using the Email Link or Copy Link functionality within BBR, the Portal generates a unique session ID for access via BBR. This ID has an internal lifespan tied to the session, which is controlled by this setting.

Session Merge Timespan Radius (Hours)
When session merge is enabled for Browser-Based Replay, you can specify the number of hours in front of and behind a session fragment in which to search for other fragments in the session.

- Session merge settings can be overridden for individual sessions in BBR. See "BBR Options" in the *IBM Tealeaf cxImpact User Manual*.

Session Merge Search Servers
When session merge is enabled for Browser-Based Replay, you can select the servers across which search server searches to find and assemble session fragments into a complete session.

- To search across all available servers at the time of replay, click the All Servers checkbox. If the list of servers changes, any newly added servers are automatically selected for searching for session fragments. Session merge for BBR is enabled at the user group level.
- See "CX User Administration" on page 165.

- See "cxReveal User Administration" in the *IBM Tealeaf cxReveal Administration Manual*.
- See “cxView User Administration” on page 184.
- Session merge settings can be overridden for individual sessions in BBR. See "BBR Options" in the *IBM Tealeaf cxImpact User Manual*.

Proxy Server Authentication for Replay Server

If a proxy server is used to connect the Replay Server to the Internet, you may need to configure authentication credentials for the Replay Server. For example, if BBR must pull images from the live site during playback, enabling this setting tells BBR to use the supplied credentials.

After it is enabled, please provide the username, password, and domain of the proxy server in the provided textbox. These credentials are stored by Search Server and are transparent to the end user.

- See "Configuring the Search Server" in the *IBM Tealeaf CX Configuration Manual*.

On-Demand Privacy

Through Search Server, you can configure privacy rules that are applied only during replay. Using this method of privacy enables the retention of sensitive data in the dataset for search and reporting purposes while safeguarding it from review during replay.

- See "On-Demand Privacy" in the *IBM Tealeaf CX Configuration Manual*.

Data Service Server

The Data Service Server manages data connections between Tealeaf servers and components with the Tealeaf databases. Through the Portal, you can configure the server that hosts the Data Service.

- See "Configuring the Tealeaf Data Service" in the *IBM Tealeaf CX Configuration Manual*.

  Note: In your Tealeaf deployment, you need only one instance of the Data Service Server.

Property
: **Value**

`Display Name`
: The name you want displayed under the Server heading on the IBM Tealeaf CX Server Configuration page.

`Host Name`
: The DNS name or IP address of the Data Service Server.

  Note: This value cannot be modified through the Portal after installation. For more information, please contact Tealeaf http://support.tealeaf.com.

`Port` The port number to which the Data Service Server listens.

- The default value is `23000`.

  Note: This value must be synchronized with the value used by each Server to connect to the Data Service Server. See "Configuring the Search Server" in the *IBM Tealeaf CX Configuration Manual*.

Visitor Report Server

If you have licensed IBM Tealeaf cxResults, this server identifies the server and port number that the Portal uses to communicate with the Query Server service, which in turn queries the IBM Tealeaf cxResults Visitor database.

Note: IBM Tealeaf cxResults is a separately licensable component of the IBM Tealeaf CX platform. IBM Tealeaf cxResults is no longer available as a newly licensed product as of Release 8.7. Customers that licensed IBM Tealeaf cxResults in Release 8.6 and earlier may continue to use and receive support for the product in Release 8.7 and later. For more information, please contact Tealeaf Customer Support .

Through the Tealeaf Management System, you must configure the location of the Visitor database. The Query Server references these registry settings in order to query the database.

- These settings are available in the Database tab of the shared configuration information beneath the TMS node. See “TMS WorldView Tab” on page 262.
- For more information on IBM Tealeaf cxResults, see "cxResults User Manual" in the *IBM Tealeaf cxResults User Manual*.

Report Search Server

The Report Search Server is used for distributing searches to various servers throughout the Tealeaf system.

- Since the Report Search Server is connected to the NT domain controller, it is used for authenticating valid users when NT authentication is used.

The Report Search Server configuration consists of the following properties:

Property
: **Value**

`Display Name`
: The label used to identify the server.

`Host Name`
: The DNS name or IP address of the Report Search Server. It must be resolvable by the server hosting the Portal.

  Note: Do not use `localhost` for the Host Name value. If the Report Server ever hosts any databases, some reporting features can be corrupted.

`Alternate URL for Link Generation`
: If users are unable to directly connect to the Portal's hostname, you enter a value here to provide user access. An example alternate URL should include both the hostname and the web application name. For example:

  ```
  http://alternate-server/Portal
  ```

  Note: This method is useful if the Report Server and Portal application are hosted on different servers. After setting this value, an IIS reset is required.

Canister Server

A Canister Server is used for storing captured hits for active and completed sessions.

As a visitor explores your site, hit after hit is captured and added to visitor's session stored in the Short Term Canister through the Canister Server. When the session has completed, the Canister Server moves the session to the Long Term Canister.

- When an instance of this server is created, an instance of the Canister Search Server is automatically created. The Portal needs to be aware of the Search Server hosted by the Canister to query for retrieval of sessions from the Canister.

Performance of the Tealeaf system may be improved by adding a new Canister Server or changing properties in your current server.

Note: All Canisters listed as `Active` in the Portal Management page must be receiving data. If a single Canister is unavailable due to network issues or other reasons, scorecards, dashboards, and other reports are not updated until the problem is corrected. When the problem is corrected, any spooled data must be collected, aggregated, and stored in the database before reports are updated with the latest data.

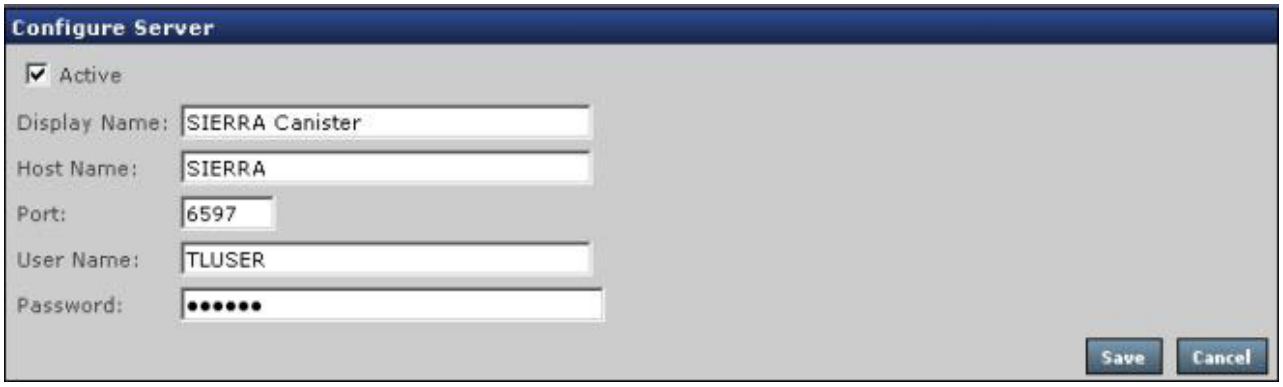

Figure 56. Canister Server

In a typical installation, the machine hosting the Canister Server may also include a Canister Search Server.

- See “Canister Search Server” on page 133.

The following fields apply to the Canister Server configuration:

Field **Value**

`Active`
When set to Yes, this value enables the data collector to retrieve data from the specified server. Set it to No to make the server definition inactive so the data collector does not retrieve data from the specified server. This field is displayed on the IBM Tealeaf CX Server Page.

Note: Reporting Data from inactive Canisters is still trimmed by the Data Collector based on the date. To preserve reporting data from Canisters which are now inactive, you must archive it to a separate location.

`Display Name`
The name you want displayed under the Server heading on the IBM Tealeaf CX Server Configuration page.

Host Name
Enter the actual host name or IP address of the computer hosting the IBM Tealeaf CX datastore. This computer collects statistics, report data, and search.

Note: The entered host name cannot resolve to the host name, IP address, or full qualified domain name of another active or inactive Canister Server.

Port The port number for the ODBC connection made to the IBM Tealeaf CX datastore.

Note: This setting is not used. The port number used is fixed as `5597`.

User Name
The user name required to access the c-tree server.

Password
The password required to access the c-tree server.

Note: Changing the default password for this server is not supported in this version.

Canister Search Server

The Canister Search Server manages queries for session data across the Tealeaf environment and other search-related functions.

- See "Configuring the Search Server" in the *IBM Tealeaf CX Configuration Manual*.

Property
Value

Display Name
The name you want displayed under the Server heading on the IBM Tealeaf CX Server Configuration page.

Host Name
The DNS name or IP address of the Canister Search Server.

Port The port number to which the Canister Search Server listens.

- The default value is `19000`.

Note: This value must be synchronized with the value specified in Canister Search Server configuration for the Canister Search Server port. See "Configuring the Search Server" in the *IBM Tealeaf CX Configuration Manual*.

Selective Archive Server

Note: cxVerify is a separately licensed component of the CX platform. This option is present only if you have licensed cxVerify prior to upgrading to Release 8.7.

A Selective Archive Server is used by cxVerify to store extracted sessions. For more information, see "Create Archive Server in the Portal" in the *IBM Tealeaf cxVerify Administration Manual*.

Note: cxVerify is a separately licensed component of the CX platform. For more information, please contact your Tealeaf representative.

Property
Value

Display Name
The name you want displayed under the Server heading on the Tealeaf CX Server Configuration page.

Host Name
The DNS name or IP address of the Archive Server.

Port The port number to which the Archive Server listens.

- The default value is 19101.

Transport Server

This section is used for configuring Transport (aka HBR) server, enabling the Portal to report status for those servers. The Transport Server configuration consists of the following properties:

- See "Health-Based Routing (HBR) Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

Property
Value

Display Name
The label used to identify the server.

Host Name
The DNS name or IP address of the Transport Server. This value must be able to be resolved by the server that hosts the Portal.

Port The port used to communicate between the Portal application and the Transport server.

TLI Server

A TLI server is used for managing static archives, which are daily and monthly archives of static content referenced in the session data stream. These archives are automatically updated and available for use during session replay.

- A .TLI file holds a static archive. A **static archive** contains content referenced in session data that is repeated frequently and does not change often. Typically, this content is considered uninteresting. For archival purposes, it may be an enterprise requirement capture this content for long-term storage in the data warehouse, which ensures the highest quality replay at any point in the future.
- For more information on static archives, see “Managing Static Archives” on page 374. The reference to the TLI Server in the Portal enables the Portal application to extract status information from the TLI Server.

  Note: Create only one (1) TLI server. Multiple TLI servers in the Tealeaf environment is not supported at this time.
- This extract information includes the list of TLI files currently maintained on the server. See “TLI Files” on page 135.
- TLI jobs to create and trim TLI files must be enabled through the Scheduling Service. See "Configuring the Scheduling Service" in the *IBM Tealeaf CX Configuration Manual*.

Property
Value

Active When set to Yes, the TLI server is active.

Display Name
The label used to identify the server.

Host Name
The DNS name or IP address of the TLI server.

Note: This value must correspond to one of the servers available through TMS and must be resolvable by the server that hosts the Portal. See Chapter 4, "Tealeaf Management System," on page 259.

Port The port of the Web server on the server.

Note: This value should match the value specified for the Search Server on the physical host. By default, Search Server uses port 19000.

TLI Files

Through the Portal Management page, you can query the TLI server to display all TLI files currently stored on the server.

- To see the list of TLI files, click **TLI Files**.

☐ List Backup TLIs

File Name	Objects	Size	FirstUse	LastUse
StaticFiles_20110902 - Copy.tli	322	2,097,152	09/02/2011 07:26:56	09/02/2011 07:26:56
StaticFiles_20110902.tli	322	2,097,152	09/02/2011 07:26:56	09/02/2011 07:26:56
MergedStaticFiles_20110930 - Copy.tli	381	3,145,728	08/31/2011 07:06:50	08/31/2011 07:06:50
MergedStaticFiles_20110930.tli	381	3,145,728	08/31/2011 07:06:50	08/31/2011 07:06:50
MergedStaticFiles_20110831 - Copy.tli	14,756	31,457,280	08/04/2011 22:08:10	08/04/2011 22:08:10
MergedStaticFiles_20110831.tli	14,756	31,457,280	08/04/2011 22:08:10	08/04/2011 22:08:10
StaticFiles_20110826 - Copy.tli	1,482	6,291,456	08/26/2011 13:03:16	08/26/2011 13:03:16
StaticFiles_20110826.tli	1,482	6,291,456	08/26/2011 13:03:16	08/26/2011 13:03:16
StaticFiles_20110825 - Copy.tli	135	1,048,576	08/26/2011 00:02:13	08/26/2011 00:02:13
StaticFiles_20110825.tli	135	1,048,576	08/26/2011 00:02:13	08/26/2011 00:02:13
StaticFiles_20110820 - Copy.tli	21	1,048,576	08/20/2011 12:18:01	08/20/2011 12:18:01
StaticFiles_20110820.tli	21	1,048,576	08/20/2011 12:18:01	08/20/2011 12:18:01
StaticFiles_20110819 - Copy.tli	767	3,145,728	08/19/2011 14:08:25	08/19/2011 14:08:25
StaticFiles_20110819.tli	767	3,145,728	08/19/2011 14:08:25	08/19/2011 14:08:25
StaticFiles_20110818 - Copy.tli	2,286	7,340,032	08/18/2011 14:02:07	08/18/2011 14:02:07
StaticFiles_20110818.tli	2,286	7,340,032	08/18/2011 14:02:07	08/18/2011 14:02:07
MergedStaticFiles_20110731 - Copy.tli	11,587	26,214,400	06/30/2011 07:09:09	06/30/2011 07:09:09
MergedStaticFiles_20110731.tli	11,587	26,214,400	06/30/2011 07:09:09	06/30/2011 07:09:09
MergedStaticFiles_20110630 - Copy.tli	12,500	29,360,128	05/31/2011 12:43:05	05/31/2011 12:43:05
MergedStaticFiles_20110630.tli	12,500	29,360,128	05/31/2011 12:43:05	05/31/2011 12:43:05
StaticFiles_20110624 - Copy.tli	303	2,097,152	12/31/1969 16:00:00	12/31/1969 16:00:00
StaticFiles_20110624.tli	303	2,097,152	12/31/1969 16:00:00	12/31/1969 16:00:00
MergedStaticFiles_20110531 - Copy.tli	45,764	85,983,232	04/30/2011 15:40:44	04/30/2011 15:40:44
MergedStaticFiles_20110531.tli	45,764	85,983,232	04/30/2011 15:40:44	04/30/2011 15:40:44

1 Page **1** of **1** (24 items)

Figure 57. TLI Files

Note: This list includes daily and monthly .TLI files.

- To view the backup TLI files on the server, select the List Backup TLIs checkbox.

- For more information on creating backups, see “Managing Static Archives” on page 374.

Field **Description**

Filename
The filename of a TLI on the TLI server.

Objects
The number of static objects in the TLI file

Size The size of the static file in B

FirstUse
The timestamp when the file was created and first used
- Use values are not populated for backups.

LastUse
The timestamp when the file was most recently used
- Use values are not populated for backups.

Note: For merged TLI files, `FirstUse` and `LastUse` values are updated on an ongoing basis.

Filenames for TLI files can be in the following formats:

Filename Pattern
Description

StaticFiles_YYYYMMDD.tli
Filename for a daily TLI file.
- `YYYYMMDD` corresponds to the date of capture and insertion of the static content

MergedStaticFiles_YYYYMMDD.tli
Filename for a TLI file merged from successive daily TLI files.
- `YYYYMMDD` corresponds to the date of last daily capture of static content for the month.

Through RTV, you can examine the contents of individual TLI files. See "Using Static Archives in RTV" in the *IBM Tealeaf RealiTea Viewer User Manual*.

cxConnect Server

This section is used for configuring IBM Tealeaf cxConnect for Data Analysis servers. See "cxConnect for Data Analysis Administration Manual" in the *IBM Tealeaf cxConnect for Data Analysis Administration Manual*.

Note: IBM Tealeaf cxConnect for Data Analysis is a separately licensed component of the IBM Tealeaf CX platform. please contact your IBM Tealeaf representative.

The IBM Tealeaf cxConnect for Data Analysis Server configuration consists of the following properties:

Property
Value

Display Name
The label used to identify the server.

Host Name
The DNS name or IP address of the IBM Tealeaf cxConnect for Data Analysis server. This value must be able to be resolved by the server that hosts the Portal.

Port The port of the Search Server that listens to the IBM Tealeaf cxConnect for Data Analysis machine.

cxVerify Server

Important: After you have upgraded cxVerify to Release 8, you may need to add back your cxVerify Server reference to the Portal Management page. After you complete the software installation, add the server reference through the Portal. This known issue may be fixed in a later build.

This section is used for configuring IBM Tealeaf cxVerify servers. See "cxVerify Administration Manual" in the *IBM Tealeaf cxVerify Administration Manual*.

Note: cxVerify is a separately licensed component of the CX platform. This option is present only if you have licensed cxVerify prior to upgrading to Release 8.7.

The IBM Tealeaf cxVerify Server configuration consists of the following properties:

Property
Value

Display Name
The label used to identify the server.

Host Name
The DNS name or IP address of the IBM Tealeaf cxVerify server. This value must be able to be resolved by the server that hosts the Portal.

Port The port of the Search Server that listens to the IBM Tealeaf cxVerify machine.

Capture Application Server

This section is used for configuring IBM Tealeaf CX Passive Capture Application servers for the Portal.

Note: To include statistics from the IBM Tealeaf CX Passive Capture Application in the Tealeaf Status report, you must create an instance of the Capture Application Server from which the Portal can retrieve statistics. See “Tealeaf Status Report” on page 231.

Note: Additional configuration of the IBM Tealeaf CX Passive Capture Application itself is available through the PCA web console. See "Passive Capture Configuration via Web Console" in the *IBM Tealeaf Passive Capture Application Manual*.

The Capture Application Server configuration consists of the following properties:

Property
Value

Use HTTPS
When selected, the server uses HTTPS only.

Note: If this value is selected, please verify that the Port number is configured appropriately.

Display Name
The label used to identify the server.

Host Name
The DNS name or IP address of the IBM Tealeaf CX Passive Capture Application server. This value must be able to be resolved by the server that hosts the Portal.

Port The port of the Web server on the IBM Tealeaf CX Passive Capture Application server.

Username
If configured, enter the username for the Portal to use to access the PCA web console.

- See "Passive Capture Configuration via Web Console" in the *IBM Tealeaf Passive Capture Application Manual*.

Password
If configured, enter the password for the Portal to use to access the PCA web console.

Note: Changing the default password for this server is not supported in this version.

- See "Passive Capture Configuration via Web Console" in the *IBM Tealeaf Passive Capture Application Manual*.

When a Capture Application server is selected, the following commands are available:

- **Ping Server** - Returns statistical information `stats.xml` from the Capture Application Server.
- **Open Console** - Opens the PCA console using the username and password specified for the Capture Application Server.

Configuring Browser-Based Replay

Through the Portal, you can configure **Browser Based Replay**, which enables the replay of captured sessions through the browsers of Tealeaf users. Browser-Based Replay (BBR) consists of two components:

1. **Replay Server** is a Windows service. The Replay Server is an HTTP server enhanced with several commands for opening sessions and preparing their pages for replay. The Replay Server typically lives on the server designated as the Report Server, which is typically the same server that hosts the Portal.
   - Using Search Server, the Replay Server is able to retrieve sessions from multiple processing servers. In a Tealeaf deployment, there must be at least one Replay Server.
2. **Portal UI** is part of the Portal. A replay user interacts with the Portal UI, which then makes requests to the Replay Server.
   - For more information on the Portal user interface, see "Browser Based Replay Interface" in the *IBM Tealeaf cxImpact User Manual*.

How BBR Works

A typical sequence of BBR operations works as follows:

1. Through the Portal UI, a Tealeaf user makes a request to open a session.
2. Replay Server prepares a query and submits it to Search Server to look for the session in the canister. If the requested session exists, Replay Server retrieves the data and returns a handle to the Portal UI, which is used for subsequent calls.
3. The Replay Server starts a thread to examine the pages in the session and to build the NavList.
   - The **NavList** is the list of viewable pages that the user can replay.
   - When the Portal UI makes a request to obtain the NavList. the returned information indicates a completion percentage for building the NavList. The Portal UI continues to make requests until the NavList is complete.
4. Based on the NavList, the Portal UI can start requesting specific pages, generally starting with the first page in the NavList.
5. When the Replay Server gets a request for a page, it prepares the page for replay and delivers it for display.
6. The Portal UI displays the page.
7. The Portal UI continues to request pages based on user actions.
8. When user review is complete, the Portal UI tells the Replay Server to close the session, which frees up all required memory.

How BBR accesses images

When a session is replayed, it often contains references to static content, such as stylesheets, JavaScripts, and images. Since images are repetitive and can be large in terms of data storage, Tealeaf provides multiple mechanisms for managing image content, so that BBR minimizes the transfer of data across the network and the impact on the server of origin.

When BBR detects a reference to an image file, the following sources are checked for a current copy of the image, in the order listed below.

1. **Internal splicing.** Depending on Replay Server and BBR settings, there may be replay-specific replacements made to referenced images.
   - For example, when `FakeImagesonRender` is enabled, the Replay Server generates a stub image for each requested image file, which reduces the amount of data across the network.
   - See "Configuring the Replay Server" in the *IBM Tealeaf CX Configuration Manual*.
2. **Replay Server cache.** The Replay Server maintains a cache of content that is from external third-party servers. This cache is queried for a reference to the requested image.
3. **BBR cache.** The Internet Explorer instance inside of BBR retains a cache, which may be used to retrieve content.
4. **TLI server.** If a TLI server has been deployed in your environment, BBR queries it for image content.
   - See “Managing Static Archives” on page 374.
5. **Origin server.** If all of the above steps fail to retrieve the image, BBR queries the server of origin.

Replay Configuration

Browser-Based Replay can be configured through the Replay Server configuration in the Tealeaf Portal.

See “Managing Tealeaf Servers” on page 115.

Replay Rules

Replay rules are used to modify a session for viewing in BBR.

Because BBR is executed within the browser, the replayed session can have more impact on the browser window than in RTV. The replay rules are mapped to a domain, and the rules for that domain are applied to any replayed sessions that match.

- See "BBR Replay Rules" in the *IBM Tealeaf cxImpact User Manual*.

Replay Rules Text File

On the master Replay Server, you can optionally edit the replay rules configuration file directly.

See "BBR Replay Rules" in the *IBM Tealeaf cxImpact User Manual*.

On-Demand Privacy

You may define and apply a set of privacy rules to data delivered from Replay Server to BBR users. These rules are useful for blocking or encrypting data that you wish to retain in the session but do not wish to expose through replay. They are applied to BBR only by the Search Server at the time of request from a Canister or archive.

- See "On-Demand Privacy" in the *IBM Tealeaf CX Configuration Manual*.
- These rules are managed and applied through the Search Server. See "Configuring the Search Server" in the *IBM Tealeaf CX Configuration Manual*.

Authentication

The Replay Server uses the same authentication mechanism as the Search Server. All requests sent to the Replay Server must have the same user information in their authentication headers as requests sent to Search Server.

- The Replay Server gets its authentication configuration from the Search Server used by machine hosting the Replay Server. See "Configuring the Search Server" in the *IBM Tealeaf CX Configuration Manual*.

Authentication Issues with Replay

If your Tealeaf system is operating in an authenticated environment, you may experience some issues during replay. If replay references some files stored on an authentication server whose permissions are not included in Tealeaf user accounts, Tealeaf users cannot retrieve those files and may not be able to experience the session as the visitor created it.

For example, suppose some Javascript files are stored behind an authentication server and Tealeaf users do not have access to this server. During replay, any references to these pages fail to retrieve and execute the Javascripts, which can result in many different kinds of issues in replay.

- If the file behind the authentication server is an image file, then the image is missing from the screen in the browser-based or RTV replay.
- By default, Tealeaf does not retain static reference content such as images, stylesheets, and Javascripts, as these items require significant additional storage and rarely change. As a result, Javascripts that are blocked from Tealeaf users cannot be applied during replay, and the effects on replay depend on the functionality of the Javascript.

To correct the replay issue with authentication, you can make one of the following changes:

1. **Move the referenced files outside of the authentication server.** You can place them somewhere within your corporate firewall in a location that is accessible to Tealeaf accounts.
2. **Copy the referenced files outside the authentication server.** If you do not wish to make changes to the web application, you can create a copy of the referenced files in an accessible location. From within Tealeaf replay, you can create rules to map to this new location.

   Note: Each time that these files are updated on the live web application, you must make the updates available to Tealeaf when session data begins to reference and require them.
   - See "BBR Replay Rules" in the *IBM Tealeaf cxImpact User Manual*.

Other Supported Formats

- **PDF** - The Replay Server displays PDFs in the replay window. No special handling by the Portal UI is required.
- **Custom Javascript Insertion** - The Replay Server inserts Javascript files near the beginning and end of each HTML page to allow for additional customization.
- **Highlighting** - The highlighting of fields and links is managed by the Portal UI.
- **Client user interface events** - If UI Capture is deployed in your web application, BBR supports limited playback of client user interface events. See "Tracking Interactions through BBR" in the *IBM Tealeaf cxImpact User Manual*.

Configuring Search Templates

Through the Portal, you can define search templates and session list templates.

- **Search templates** specify the set of fields and default properties that are available to users when the search template is selected. See “Search Templates.”
- **Session list templates** are lists of attribute columns used on the session list grid for search results. See “Session List Templates” on page 155.

To define search templates, select **Configure** > **Search Templates** in the Portal menu.

For more information about configuring search templates, see "Configuring a Search Template" in *IBM Tealeaf cxReveal Administration Manual*.

Search Templates

A search template defines the scope of the criteria that can be used to search and are an effective means of limiting for users and groups the searchable fields that they can search.

By developing and deploying well-designed templates, the Tealeaf administrator can focus searches on the terms of interest to specific groups, prevent access to sensitive fields, and limit the number of returned results for better search performance.

- Search templates are available for selection in the session search and visitor search pages.
  - See "Searching Session Data" in the *IBM Tealeaf cxImpact User Manual*.
  - See "Searching for Visitors" in the *IBM Tealeaf cxResults User Manual*.

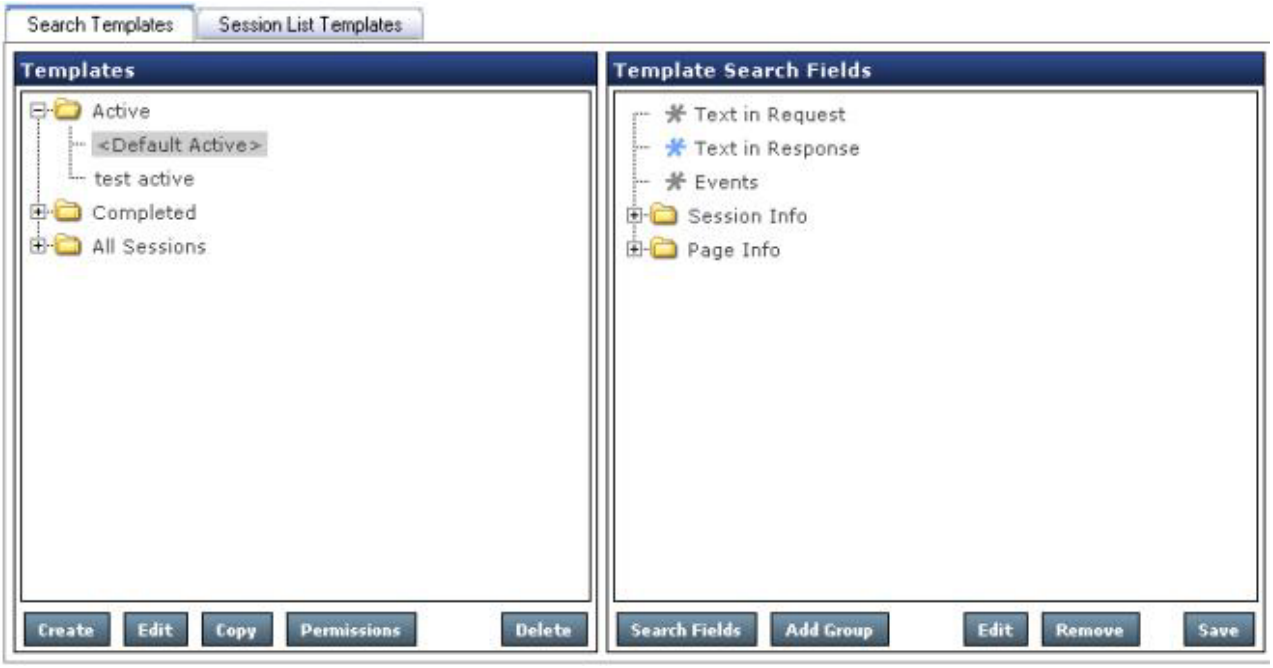

Figure 58. Search Templates

Search templates control the fields that are available for a user to use in a search, as well as the default attributes of the search such as date range and servers to search. On the Search Templates pane of the Search Template Configuration page, there is a list of defined templates, categorized by search type. The three search types correspond to a set of searchable sessions:

- **Active** searches the currently active sessions.
- **Completed** searches completed and indexed sessions
- **All Sessions** searches across both active and completed sessions.

As the templates are incompatible across search types, each template has a list of search fields available to it.

Configuring a Search Template

Note: After you have created a new search template, you should immediately set the permissions to enable viewing of it. Otherwise, it does not appear among your available search templates. To edit permissions, select the template and click **Permissions**. See “Permissions” on page 145.

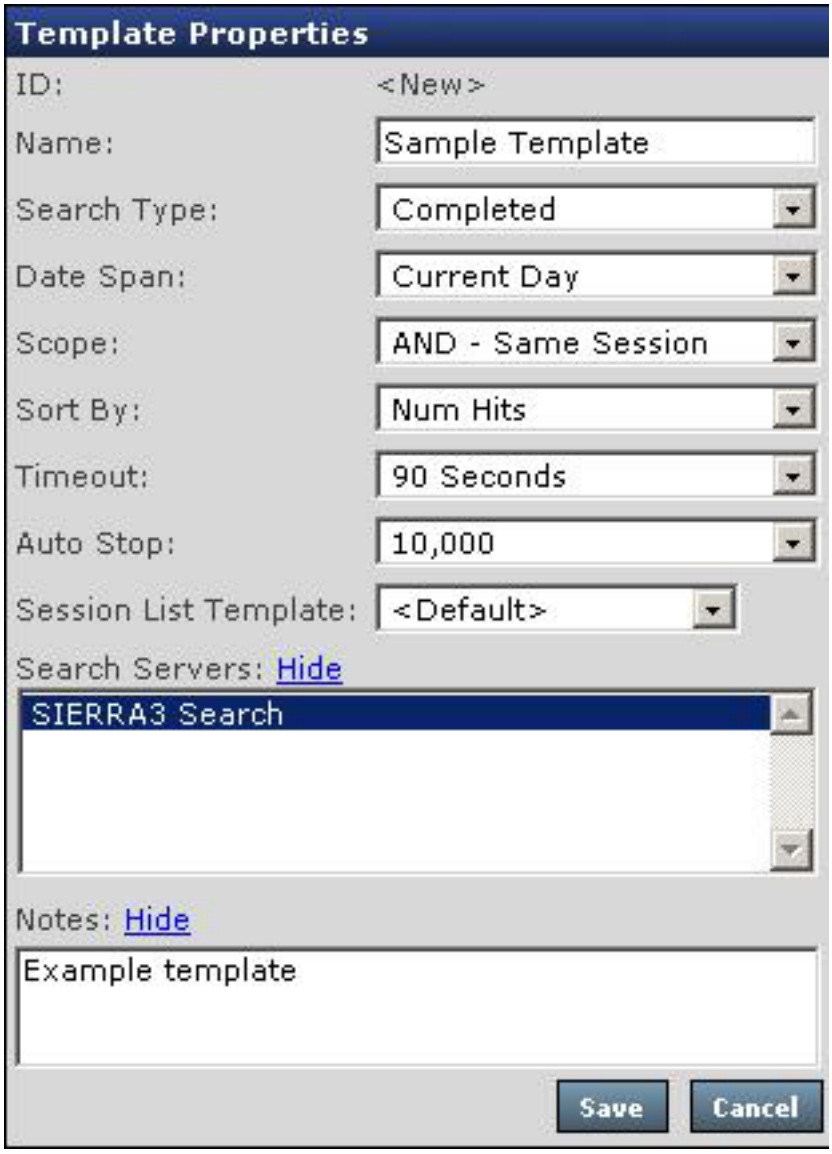

Figure 59. Template Properties

The available properties for a search template are listed below. Each property on the template is a default setting, which the user may alter at the time of the search, if enabled by permissions:

Property
Description

Name The name of the template.

Search Type
The type of sessions that can be searched using the template: `Active`, `Completed`, or `All Sessions`.

Date Span
The date range over which the search template covers.

Scope How selected fields are applied to the search:

- `AND - Same Session` - The search returns sessions that match all specified fields.

- AND - Same Page - The search returns sessions that have individual pages matching all fields.
- OR - Any Session - The search returns sessions that match any of the fields.

Sort By
How the matching sessions are sorted.

Timeout
The maximum amount of time allotted for the search.

Auto Stop
The search automatically stops after opening the number of sessions defined by this property.

Session List Template
The Session List template in which to display the results returned when a search is executed with this search template. See "Session List Templates" on page 155.

Search Servers
The search servers to select for searching by default.

Notes If specified, the notes appear as an informational textbox above the search fields on the search page.

- To insert a newline in the Notes textbox, enter \n. These text strings appear as-is in the Notes textbox. You can press ENTER in the textbox to make the notes more legible.

Copying Templates

You can copy a search template and modify it to create your own search template. To copy a template, click **Copy** beneath the list of templates.

If the Duplicate Search Fields option is selected, new copies of the search fields are created with the new search template. Otherwise, the new template uses the same fields as the original. Because search fields and templates are independent, multiple templates can use the same search field.

Permissions

You can specify the user groups that can access the selected template for searches.

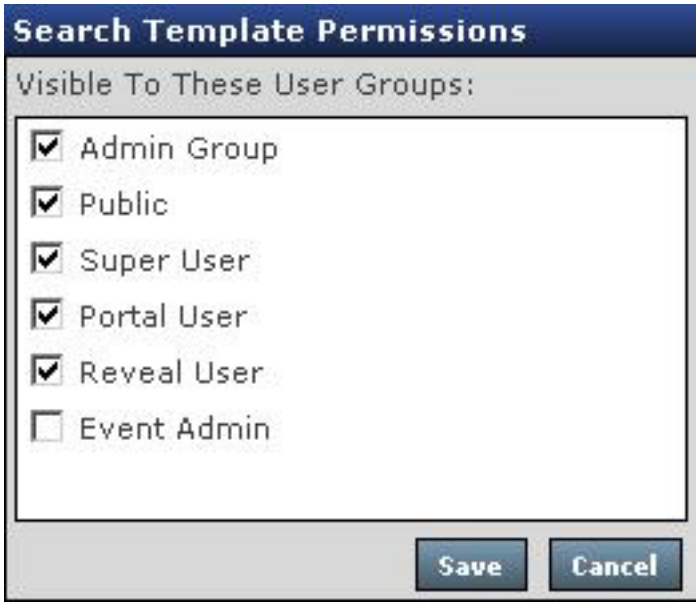

Figure 60. Search Template Permissions

Deleting Templates

To delete a template to which you have access, click **Delete**. Any search fields used by the template are not deleted.

Template Search Fields

Search fields are defined in a search and populated with values to generate a specific search query. Individual search fields can be configured and assigned to search templates

- To add a new search field group, click **Add Group**. Enter a name and select **Add**. In the Search Template Fields panel, a new folder appears with the group name as its label.
- To edit a template search field, select it and click **Edit**. See "Search field properties" on page 151.
- Search fields may be labeled as Available, Active, or Required. For more information, please review the Search Field Modes panel.
- To change the mode for a search field, click the icon next to it and select the new mode. See "Search Field Modes" on page 153.

  Note: A search template can display up to six recommended or required fields. These fields appear when the template is first loaded. If additional recommended or required fields are specified in the template, only the first six such fields are displayed when the template is initially displayed to the user.
- To remove a field, select it and click **Remove**.
- To save the search template, click **Save**.

Configuring a Search Field

To add a search field to the template, click **Search Fields** in the Template Search Fields pane to display the currently available search fields:

Note: If a field does not appear in the list, it may be incompatible with the current template and is therefore hidden. Incompatible fields use search keywords that cannot be used with the current template. To display all fields, select the **Show Incompatible Fields** checkbox. Incompatible fields cannot be added to the current template.

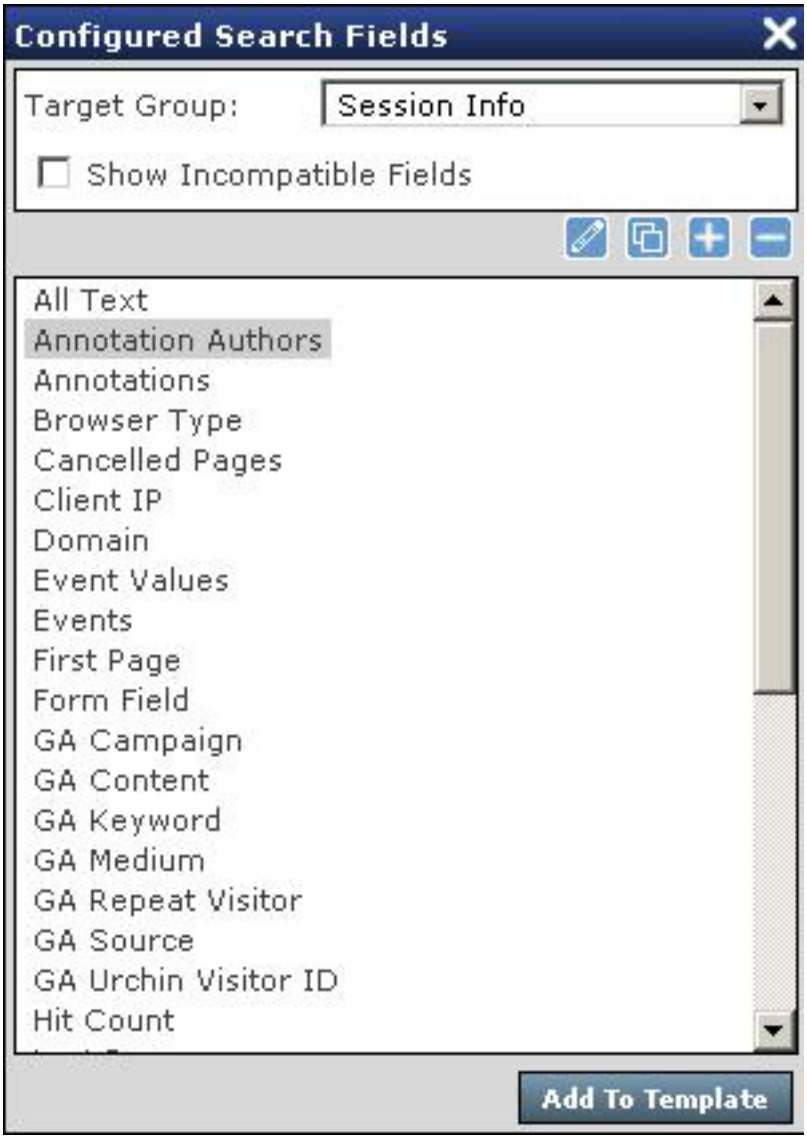

Figure 61. Configured Search Fields

- To add a field to a field group, select the group from the Target Group drop-down.
- To add a field to the current search template, select it and click **Add To Template**.
- To edit a search field, click the Edit icon. See "Search field properties" on page 151.
- To create a new search field, click the + icon. See "Creating a New Search Field" on page 150.

Available Search Fields

The following search fields are available in the default templates for Active and Completed sessions.

Note: Depending on the type of search template, some search fields may not be available. For All Sessions search templates, you may add search fields only if they are applicable to both active and completed search templates.

Note: Fields that have been encrypted using privacy rules in the IBM Tealeaf CX Passive Capture Application or Windows pipelines cannot be decrypted in the Portal.

- These encrypted fields can be decrypted **only** during replay.
- As an alternative, you can leave the configured fields in unencrypted state in the session data and then define privacy rules specifically to be applied during session replay, permitting the display of the unencrypted data in the Portal, as needed. See "On-Demand Privacy" in the *IBM Tealeaf CX Configuration Manual*.

Text:

For completed sessions, you can search for text strings contained in requests, responses, and other derived data. The text strings can include different search patterns to further restrict the search.

Note: Free text searches can be expensive to complete. These search fields should be included only where other indexed search field options are not available.

Text String Type
Description

`All Text`
When searching completed sessions, you can search for a text string in any part of the session data, including the request, the response, and any derived data such as events or custom user-defined fields. This data is indexed using the dtSearch software program, so any of the search options provided by dtSearch are available. The following patterns are supported:

- `*` Wild character. For example, `C*` matches both `CC` and `CAT`.
- `{ }` Matches any character contained within the bracket.
- `^` Represents the beginning of the line. For example, `^T` searches for any line starting with a `T`.
- `$` Represents the end of the line. For example, `error$` returns any line ending with `error`.
- `\` indicates to emphasize the subsequent character in the search.

`Text in Request`
You can search for a text string in the request portion of a hit in completed sessions. The above patterns are supported.

`Text in Response`
You can search for a text string in the response portion of a hit in completed sessions. The above patterns are supported.

Events:

During processing of sessions, Tealeaf events are triggered based on visitor session data. The data derived from events is stored with the actual hit and session information and may be searched like captured data.

- With IBM Tealeaf cxResults, these events are also stored in the Visitor database where they can be searched for dates beyond the life of the replayable sessions.

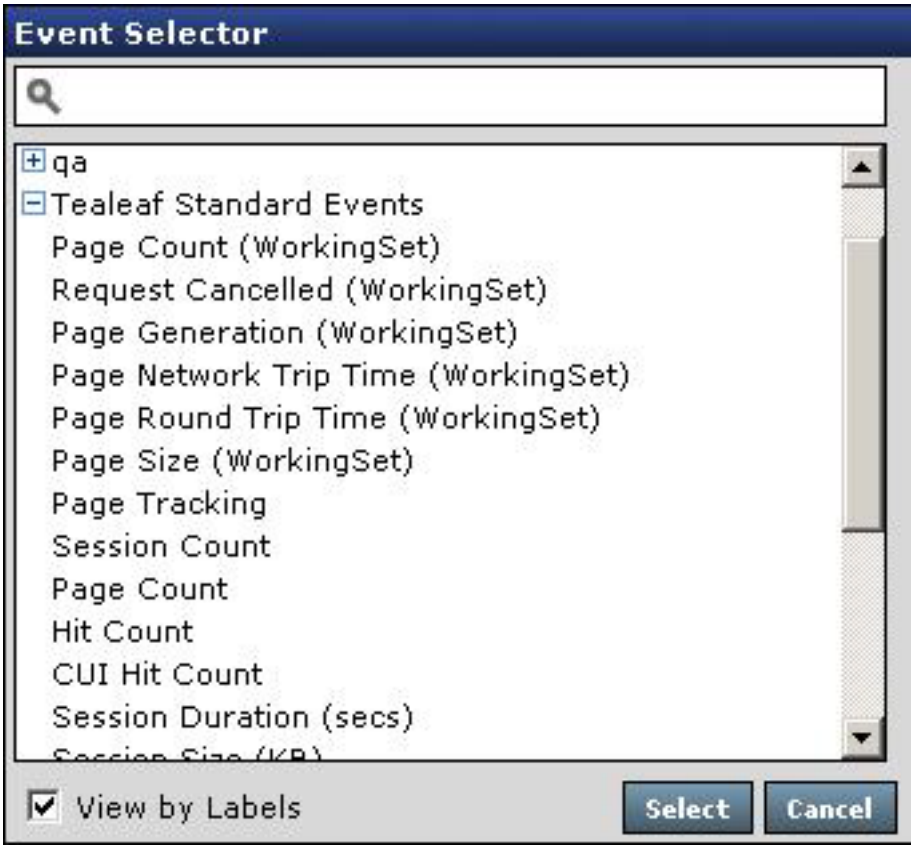

Figure 62. Event Selector

Events may be associated with one or more dimensions, which enable the reporting of event data across multiple axes.

- See "Tealeaf Data Model" in the *IBM Tealeaf Reporting Guide*.
- Users of the search template may select the dimension or dimensions associated with the event to search as part of the query. See "Searching Session Data" in the *IBM Tealeaf cxImpact User Manual*.
- Dimensions are created in the Tealeaf Event Manager. See "TEM Dimensions Tab" in the *IBM Tealeaf Event Manager Manual*.

Input Type for events

Search fields for events may be configured to use one of the following input types:

Input Types
Description

`Event Selector`
Template user may select any accessible event.

`Pulldown`
Template user may select an event from a pulldown menu that you populate. You may also specify default values.

`Radio buttons`
Template user may select an event by clicking a radio button from a set that you define. You may also specify default values.

Event Values:

You may specify fields that enable search for specific event values. Users may select the event and value for which to search or use the default ones you specify.

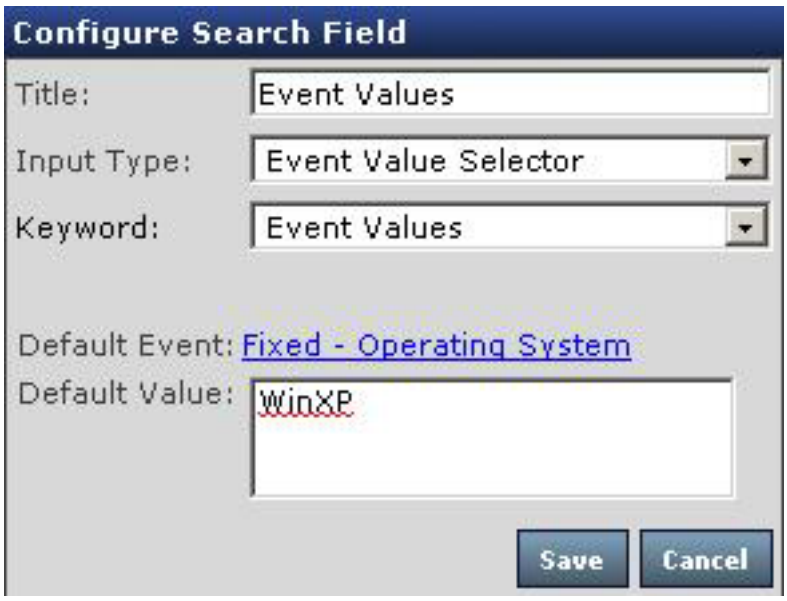

Figure 63. Configuring Event Value Search Field

Field Description

`Default Event`
Click the link to select the event to use as the default event for this search field.

`Default Value`
Enter the default value with which to populate the search term. No data validation of the default value is performed.

Event values may be specified using different kinds of input types. See “Input Type for events” on page 148.

Values for events are stored in a data entity called a **fact** in the request buffer.

- For more information on session attributes, see "RealiTea Viewer - Request View" in the *IBM Tealeaf RealiTea Viewer User Manual*.

Session Info:

As Tealeaf builds and processes sessions, it derives data from the captured information that is applicable to the entire session.

For more information on the default fields provided in active and standard search templates, see "Searching Session Data" in the *IBM Tealeaf cxImpact User Manual*.

Page Info:

As Tealeaf creates and processes hits, it derives data from the captured information that applies to the entire hit or page.

- For this grouping, the list of searchable fields for active sessions is completely different from completed sessions.
- The active page fields are aggregated into reports contained within dashboards.

For more information on the default fields provided in active and standard search templates, see "Searching Session Data" in the *IBM Tealeaf cxImpact User Manual*.

cxReveal Fields:

Tealeaf enables IBM Tealeaf cxReveal users to search for sessions by session attribute. These session attributes are stored in a separate high-performance database, where they are accessible for search as soon as they are identified in the session data.

Note: Searching for session attributes is a component of cxReveal, a separately licensable component of the Tealeaf CX platform. please contact your IBM Tealeaf representative.

When configuring search templates of IBM Tealeaf cxReveal type, you may select from the available session attributes that have been marked to be searchable and tracked in the IBM Tealeaf cxReveal database.

Note: When adding a IBM Tealeaf cxReveal search field to a search template, you should verify that it is specified with a valid session attribute that is collecting data. Open the search field definition to verify that the session attribute is not listed as `Unknown Attribute`. Such search fields can be added to a search, yet they return zero sessions every time. Each IBM Tealeaf cxReveal search field should be verified independent of other IBM Tealeaf cxReveal search fields.

- You may define the session attributes that are marked for searching and tracking in the database in the Event Manager. See "TEM Session Attributes Tab" in the *IBM Tealeaf Event Manager Manual*.

Note: These search keywords are compatible with the IBM Tealeaf cxReveal Active, Completed, and All Sessions search templates only, as these templates perform queries of the IBM Tealeaf cxReveal database.

- These search template fields cannot be used in IBM Tealeaf cxImpact search templates.
- See "Configuring Session Attribute Search" in the *IBM Tealeaf cxReveal Administration Manual*.

Creating a New Search Field

To create a new search field, click + in the Configured Search Fields dialog.

Search field properties:

Creating or editing a search field opens the Configure Search Field dialog:

Figure 64. Configure Search Field

You can modify the following properties for your new or edited search field:

Property
Description

Title The title of the search field.

Input Type
The types of input the search field can use. This value affects how the search field is displayed and the users who are allowed to use it.

- **Free Text** - The user is presented with a textbox and can freely enter text. The textbox is populated with the contents of the Default Value property.
- **Pulldown** - The user is presented with a drop-down of specified values from which to select. The default selected item is specified by the Default Value property. The items in the drop-down are controlled by the Values property.
  - You can specify empty values, if desired. If the user selects the label corresponding to an empty value, the field is not submitted as part of the search.
- **Radio Button** - The user is presented with a list of radio buttons of specified values from which to select. The default selected item is specified by the Default Value property. The items in the drop-down are controlled by the Values property.
  - You can specify empty values, if desired. If the user selects the label corresponding to an empty value, the field is not submitted as part of the search.
- **Prebuilt Query** - You can enter a pre-configured query in one of the following textboxes:

- `Query` - If the search field is selected to be `included` in the search by a user of the template, this query is used in the search.
- `Inverse` - If the search field is selected to be `not included` in the search by a user of the template, this query is used in the search. If this query is not specified and the user selects `not included`, a negated version of the `Query` value is used.

- **Form Field Input** - This value is only used with the Form Field keyword. When it is used, the user is presented with two textboxes: one for a field and one for a value. The textboxes are populated with the values of the Field and Value properties, respectively.
  - **Hide Field Input** - If this is enabled, the field input for a form field is hidden, and the default is used. This feature can be used to restrict a form field search to a specific field and to permit only an input for the value.
- **Event Selector** - This Input Type value is only used with the Events keyword. When it is used, it provides the user with the Event Selector popup when the field is selected.

`Keyword`
The actual search keyword to use when this field is selected. See "Selecting keywords."

`MD5 Hashing`
For search fields that match against text, this option hashes the value the user inputs before generating the query. This feature allows a user to enter plain text to match against an indexed field that has been MD5 hashed to clean or obfuscate data.

`Default Value`
The default value to use when the search field is selected. This value changes depending on the Input Type.

- For event and event value fields, you may select the default event for the field.

`Values`
When the Pulldown and Radio Button Input Types are selected, the Values property is enabled and contains a list of the displayed values from which the user can choose.

- To add a value, click **Add**.
- To edit a value, click **Edit**.
- To remove a value, click **Remove**. You can reorder values by using the Up and Down buttons.
- **Label** - The visible text for this option.
- **Value** - The actual value for the option that is used to generate the search query.
- **Events** - Displays the Event Selector to populate the Label and Value with an event.

Selecting keywords:

The available items correspond to search keywords used when building the actual query. For example `Event` corresponds to `tltstsuniqueid` in Completed Search Templates.

Some keywords are not compatible with the current template type. When you select an incompatible keyword, the **Keyword:** marker indicates that there is an issue. Mouse over the Question Mark (?) icon to review the issue, as indicated below:

Figure 65. Issue with Keyword

If you have an incompatible keyword, you cannot save the search field until the issue is corrected.

Selecting the following keywords populates special items in the Input Type list:

Keyword
Input Type item

Event Event selector

Form Field
Form Field Input (two text inputs, form field and value)

AppData
AppData Input (same as form field input)

See "Search Keywords" in the *IBM Tealeaf cxImpact User Manual*.

Search Field Modes

A search field added to a template has a mode: Available, Recommended, and Required.

- To select the mode for a field, click the icon next to the search field.

Mode **Description**

Available
The field is not displayed but can be used in a search. This is the default mode of a search field that is added to a template.

Active The field is displayed and populated with a pre-configured default value. An active field can be removed from a search. An empty active field will not be included in a search.

Required
The field is displayed and populated with a pre-configured default value. A required field cannot be removed from a search. A required field must contain a value or valid selection when the search is performed.

Adding Session Attributes

A **session attribute** is a session-level storage container for data that can be populated by an event. Tealeaf supports the use of up to 64 session attributes.

- From pre-Release 8.0 versions of Tealeaf, the LoginID and UserDef 1-4 user-defined values have been migrated to Session Attribute 0 and Session Attributes 1-4, respectively.
- See "TEM Session Attributes Tab" in the *IBM Tealeaf Event Manager Manual*.

Note: Depending on the trigger of its populating event, a session attribute may not be available in a search template, even if you configure it. For example, if the session attribute is populated by an event triggered by End of Session, it is possible to add the session attribute as a field to an Active search template, even though the field is not yet populated for active sessions.

To add a session attribute to a search template:

1. Create the session attribute, including populating it with values generated by an event. See "TEM Session Attributes Tab" in the *IBM Tealeaf Event Manager Manual*.
2. In the Search Templates page, select the template in which to create the session attribute search field.
3. Click **Search Fields**. The Configured Search Fields dialog appears.
4. Select the Target Group from the drop-down.
   - If your search template does not include an appropriate group, close the Configured Search Fields dialog and click **Add Group**. See “Template Search Fields” on page 145.
5. Before a search field is available for use in search templates, you must create it. Click the + icon.
   a. The Configure Search Field dialog appears.
   b. Specify a title for the field, which appears in the search template and in the search configuration page.
   c. From the Keyword drop-down, select `Session Attribute`.
   d. The Attribute selector is displayed in the Configure Search Field dialog. Click **<Select an attribute**.
   e. Select the session attribute in the displayed list. Click **Select**.
   f. Configure the other properties as needed.
   g. Click **Save**.
6. The selected session attribute is added to the search template under the selected group.
7. To save changes to the server, click **Save**.

You may also add the session attribute as a column in the session list template. See “Populating a session list template column with a session attribute” on page 162.

Adding Client User Interface Fields

Tealeaf UI Capture enables the capture of user interface events that occur in the visitor's browser. When deployed as part of your web application, UI Capture tracks user interface events on the client and submits XML-based requests to the web server, which enables those events to be captured and processed by Tealeaf.

Note: IBM Tealeaf CX UI Capture for AJAX requires the Extended Edition of the IBM Tealeaf CX platform. please contact your IBM Tealeaf representative.

- For more information on UI Capture, see "UI Capture FAQ" in the *IBM Tealeaf UI Capture for AJAX FAQ*.

When UI Capture is deployed and configured for your web application, the following values are inserted as variables in the [appdata] section of the request of last viewable page before the UI event. These variables map to search keywords that you can add to your search templates.

- These values are identified in the Windows pipeline and posted into the [appdata] section by the Tealeaf Reference session agent. See "Tealeaf Reference Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

Table 7. Adding Client User Interface Fields

Request Variable	Corresponding Search Keyword	Description
TLT_CUI_URL	Client UI URL	URL as reported from the client user interface
TLT_CUI_APPLICATION_NAME	Client UI Application	Name of the application as reported from the client user interface.

The above values may differ from the standard reference values that are inserted into the request by the Tealeaf Reference session agent, as the client user interface may contain an entire application with unique URLs within a single HTML page.

Session List Templates

Session list templates are used to configure the appearance and the data of the list of sessions returned from a search. The session list templates allow control of the displayed fields of session data, how they should be treated, and other information such as the server metrics.

Search templates can be set to use specific session list templates, or a template can be selected at the time of the search. Much like the search templates, session list templates have a set of properties, but they can also reference session list columns, which are assigned to the templates.

Figure 66. Session List Templates

- To create a new session list template, click **Create**. See “Configuring a Session List Template.”
- To edit an existing session list template, select it and click **Edit**. See “Configuring a Session List Template.”
- To copy a session list template, select it and click **Copy**. See “Copying a Session List Template” on page 158.
- To assign the session list template to one or more search templates, select it and click **Assignments**. See “Assigning a Session List Template” on page 159.
- To delete a session list template, select it and click **Delete**.

Configuring a Session List Template

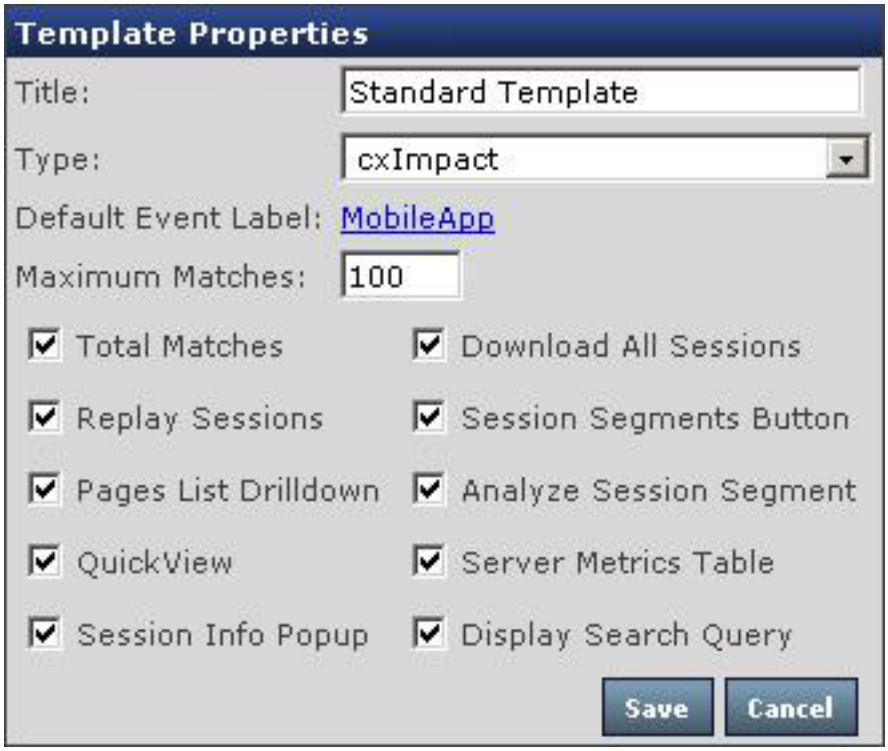

Figure 67. Template Properties Session list templates have the following properties:

Property
Description

Title The name of the template.

Type The type of search template:

- `cxImpact` - standard search template. Use this one in almost all cases.
- `cxReveal` - IBM Tealeaf cxReveal search templates can be used to specify searches by session attribute of the IBM Tealeaf cxReveal database.

  Note: The ability to search by session attribute requires IBM Tealeaf cxReveal, an installation of the IBM Tealeaf cxReveal database, and additional configuration. See "Configuring Session Attribute Search" in the *IBM Tealeaf cxReveal Administration Manual*.

Default Event Label
Select the event label whose events should be displayed by default in the QuickView page.

- Users may select a different event label as needed in the QuickView page.
- See "Searching Session Data" in the *IBM Tealeaf cxImpact User Manual*.

Maximum Matches
The maximum number of sessions to return from a search. This setting has no effect on how many sessions are found or examined; it only determines the number of sessions available to the user. For performance reasons, you cannot set this value above 10000.

- This setting should be configured in conjunction with other search limit settings. See “Configuring Maximum Session Display Limits” on page 158.

Total Matches
Displays the total number of matching sessions.

Replay Sessions
Enables the replay of sessions from the session list.

Pages List Drilldown
Allows examining the pages of a returned session within the Portal.

QuickView
Enable the user to open the QuickView page, where the events that fired on each page can be reviewed.

- See "Searching Session Data" in the *IBM Tealeaf cxImpact User Manual*.

Session Info Popup
Allows the user to get a quick list of session attributes from the session list.

Download All Sessions
Displays the Download All button, which allows users with RTV installed to open all sessions found in a search.

Session Segments Button
Displays the Manage Segments button in the session list, which allows users to access the Manage Segments page.

Note: When this button is enabled in session list templates accessible to IBM Tealeaf cxReveal users, it is not displayed in the session list because IBM Tealeaf cxReveal users cannot create session segments. IBM Tealeaf cxReveal administrators may be able to create session segments through the Manage Session Segments page.
See "Managing Session Segments" in the *IBM Tealeaf cxResults User Manual*.

Analyze Session Segment
Displays the Analyze Segment button, which can be click to analyze the session segment of a search directly from the Session List. See "Analyzing Session Segments" in the *IBM Tealeaf cxResults User Manual*.

Note: When this button is enabled in session list templates accessible to IBM Tealeaf cxReveal users, it is not displayed in the session list because IBM Tealeaf cxReveal users cannot create session segments. IBM Tealeaf cxReveal administrators may be able to create session segments through the Manage Session Segments page. See "Managing Session Segments" in the *IBM Tealeaf cxResults User Manual*.

Server Metrics Table
Displays the table listing the servers that were searched with elapsed time and number of returned sessions.

Display Search Query
Allows the query used in the search to appear on the Search Results page.

Configuring Maximum Session Display Limits:

When configuring the number of sessions that are displayed in search results, several Portal settings may affect the result.

- **Search Database Max Sessions Returned** - (IBM Tealeaf cxReveal database searches only) This parameter defines the maximum number of sessions that can be returned from a search of the IBM Tealeaf cxReveal database. For IBM Tealeaf cxReveal database searches, other session display limits may be applied to further limit the number of sessions displayed; the number of displayed search results never exceeds this number of database searches.
  - See "cxReveal Settings" in the *IBM Tealeaf cxReveal Administration Manual*.
- **Maximum Matches** - In the specified session list template, you can configure the maximum number of sessions that can possibly be displayed in the template.
  - Regardless of the values of other settings, the `Maximum Matches` setting defines the limit to be displayed in a session list.
  - For performance reasons, this value cannot exceed `10000`.
- **Session Segment Max Sessions Limit** - In the Portal Management page, you can specify the maximum number of sessions that any newly generated segment can contain.
  - See "CX Settings" in the *IBM Tealeaf cxImpact Administration Manual*.
  - See "cxReveal Settings" in the *IBM Tealeaf cxReveal Administration Manual*.
- **Live Max Results Limit or Completed Max Results Limit** - In the Portal Management page, these settings determine the number of sessions that can be displayed for any active or complete session search, regardless of the template setting.
  - See "CX Settings" in the *IBM Tealeaf cxImpact Administration Manual*.
  - See "cxReveal Settings" in the *IBM Tealeaf cxReveal Administration Manual*.
- **Auto Stop Limit** - This setting defines the number of sessions that Search Server inspects per Canister before returning results. From an individual Canister, Search Server does not inspect more sessions for matches than the number configured for this setting. By default, this value is set to `10000`.
  - See "Configuring the Search Server" in the *IBM Tealeaf CX Configuration Manual*.

The maximum number of sessions displayed in a search result is the minimum of the applicable settings. The Live or Completed limit is applicable depending on the type of search and only in cases where the search generates a new result set.

- Results can be returned into existing result sets, where the maximum limits for the existing result are applied, instead of the current Portal settings.

Limits for searches executed from IBM Tealeaf CX RealiTea Viewer can be configured in the RTV Search setup dialog. See "RTV Search Setup" in the *IBM Tealeaf RealiTea Viewer User Manual*.

Copying a Session List Template

To copy a template, select it from the list of templates and click the Copy button beneath the list.

The Duplicate Columns option controls whether or not copies are also made of the columns associated with the template. Because columns and templates are independent, multiple templates can use the same column.

- If the Duplicate Columns option is selected, new copies of the columns are created along with the new template. Otherwise, the new template uses the same columns as the original.

Assigning a Session List Template

To assign a session list template to one or more search templates, select one or more search templates and click the Assignments button beneath the list of templates. The selected search templates are updated to use the current session list template by default.

Selecting a Session List Template

In the Search page, you can select a session list template to use.

- See "Searching Session Data" in the *IBM Tealeaf cxImpact User Manual*.
- If IBM Tealeaf cxResults is licensed, you can select session list templates in the Visitor Search page. See "Searching for Visitors" in the *IBM Tealeaf cxResults User Manual*.

For completed session searches, you can choose to display them in a different session list template. See "Managing Session Segments" in the *IBM Tealeaf cxResults User Manual*.

Example Session List Templates

This section explains the use of session list templates to filter your search results.

Session counts template:

In some cases, you may only be interested in the number of instances that the search criteria occurred. You can configure a session list template to return only the counts of found results.

- When searching for session counts only, the display of the results is faster; the search itself takes the same length of time.
- Even if you are only returning the session counts in the session list, the list of returned sessions can still be exported through the Portal.

To configure:

1. Click the Session List Templates tab.
2. Create or select the session list template to use.
3. For the Maximum Matches value, enter `0`.
4. Select the Total Matches checkbox.
5. Select the Server Metrics Table checkbox.
6. The other checkboxes are irrelevant to this template.
7. Click **Save**.
   - To make this the default session list template for users or groups, you must assign the session list template. See “Assigning a Session List Template.”

Session List Columns

Session list columns are used on session list templates, and they control what is displayed in the list of matching sessions. The columns can be used to display session attributes, compare session attributes to specified values, or flag for the existence of specific events or other attributes.

- For more information on the columns in the standard session list template, see "Searching Session Data" in the *IBM Tealeaf cxImpact User Manual*.

Adding a Session List Column to a Template

Select a Session List template, and the template's columns are listed in the Template Search Columns pane. To select configured columns, click the All Columns button. To add a column to the template, select a column and click the Add To Template button. Columns in the Session List Columns pane can be reordered using drag and drop.

Configuring a Session List Column

On the Search Template Configuration page, the **All Columns** button in the Template Search Columns pane displays a popup of the currently available columns. Use the **New** button to start creating a new column, or select an existing column and use the **Edit** button.

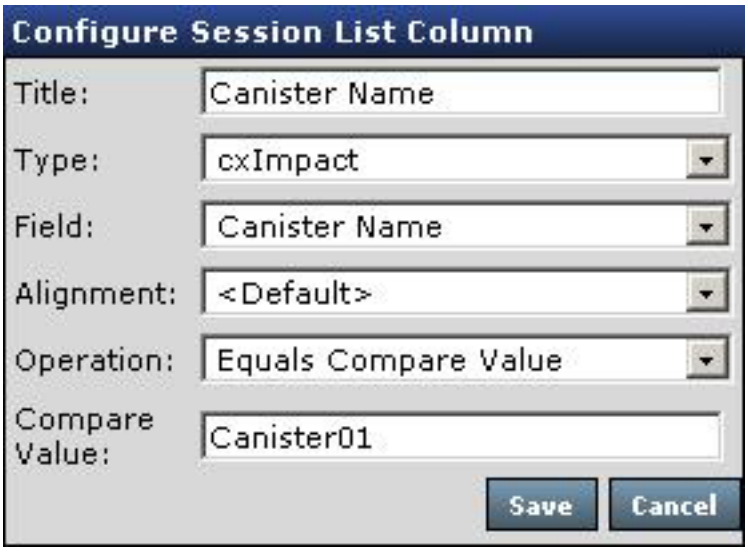

Figure 68. Configure Session List Column

For more information on configuring an event column, see “Configuring an Event Column in a Session List Template” on page 161.

Session list template columns have the following properties:

Property
Description

Title The title of the column, displayed on the session list.

Type The type of session list template column:

- `cxImpact` - standard search template column. Use this one in almost all cases.
- `cxReveal` - IBM Tealeaf cxReveal search templates can be used to specify searches by session attribute of the IBM Tealeaf cxReveal database.

  Note: The ability to search by session attribute requires IBM Tealeaf cxReveal, an installation of the IBM Tealeaf cxReveal database, and additional configuration. See "Configuring Session Attribute Search" in the *IBM Tealeaf cxReveal Administration Manual*.

Field The session field that the column uses for data.

Note: The list of available fields depends on whether standard or extended indexing is in use. For more information on indexed fields, see "RealiTea Viewer - What is Indexed and Searchable" in the *IBM Tealeaf RealiTea Viewer User Manual*.

Alignment
Optionally, you can override the default alignment setting for the specific value type to be applied to the session list column by making a selection in the drop-down.

Operation
How the data is displayed or used:

- **<Display Field Value>** - The field's value is displayed directly.
- **Equals Compare Value** - A checkmark is displayed if the value of the Field exactly matches the Compare Value.
- **Less Than Compare Value** - For numeric fields, a checkmark is displayed if the value of the Field is less than the Compare Value.
- **Greater Than Compare Value** - For numeric fields, a checkmark is displayed if the value of the Field is greater than the Compare Value.

Compare Value
The value against which to compare the Field value.

Events To List
The set of possible events to list for the Event List Subset field. The listed events are the intersection of the events of each session and the supplied list for the column. To add or remove events, click the Add or Remove buttons above the list.

Configuring an Event Column in a Session List Template:
You can configure the event icons displayed in a session list column. When an event list subset column is added to a session list template, you can provide parameters on the events for which you can search.

- In the session list, users can review details on the event. See "Searching Session Data" in the *IBM Tealeaf cxImpact User Manual*.

1. In the Session List Templates tab, select the template to edit or create a new one.
2. Select the event column to edit, or create a new one.

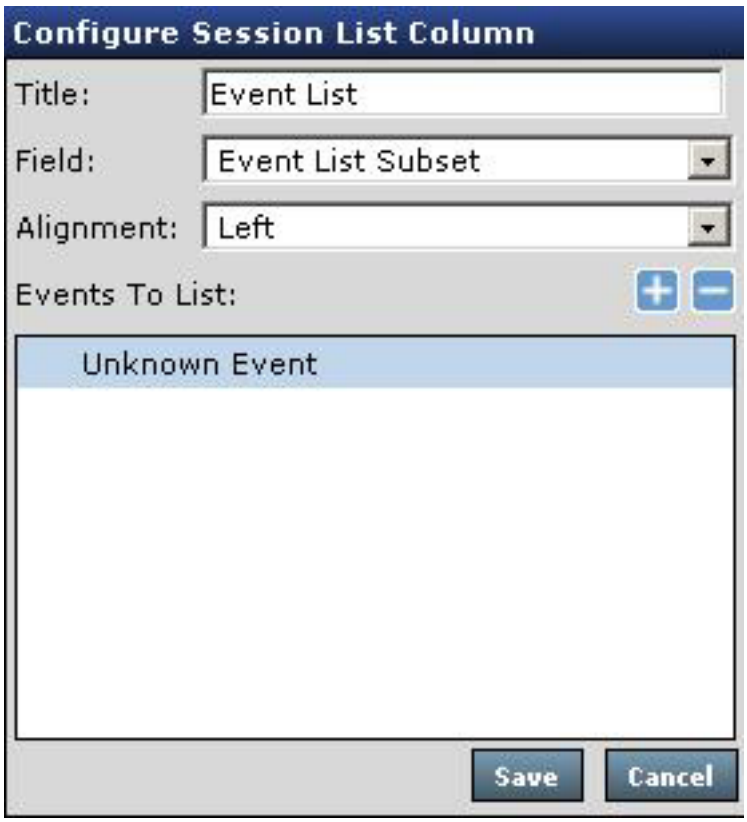

Figure 69. Configure Session List Column

3. From the Field drop-down, select `Event List Subset`.
4. Use the + and - buttons to add or remove events from the list.
5. Save the column.
6. Add the column to the template, if necessary.
7. Save the template.

Populating a session list template column with a session attribute:

Tealeaf enables the creation and tracking of up to 64 session-level parameters into distinct session attributes. If desired, you can populate session list templates with session attributes, which is useful for exporting session data.

Note: Columns populated by session attributes are sorted as text values.

- For more information on session attributes, see "TEM Session Attributes Tab" in the *IBM Tealeaf Event Manager Manual*.

To add a session attribute to a session list template:

1. Create the session attribute, including populating it with values generated by an event. See "TEM Session Attributes Tab" in the *IBM Tealeaf Event Manager Manual*.
2. In the Search Templates page, select the session list template in which to create the session attribute search field.
3. Click **All Columns**. The Configured Session List Columns dialog appears.
4. Before a session list column is available for use in search templates, you must create it. Click the + icon.

a. The Configure Session List Column dialog appears.
b. Specify a title for the column, which appears in the session list and its template.
c. From the Field drop-down, select `Session Attribute Value`.
d. The Attribute selector is displayed in the Configure Search Column dialog. Click **<Select an attribute**.
e. Select the session attribute in the displayed list. Click **Select**.
f. Select the desired Operation from the drop-down.
   - To export the values of the session attribute, select `<Display Field Value>`.
g. Click **Save**.

5. The selected session attribute is added to the list of available columns.
6. To add the new column to your template, select it and click **Add to Template**.
7. To save changes to the server, click **Save**.
   - For more information on exporting data through the session list template, see "Searching Session Data" in the *IBM Tealeaf cxImpact User Manual*.

You may also add session attributes as fields to your search template. See "Adding Session Attributes" on page 153.

Building Searches via Encoded URL

This section describes how you can build an encoded URL to execute a specified search.

Note: The best method for building searches via encoded URL is to create and execute the search through the Tealeaf Portal. When the desired results are displayed, you may step back to the search page and click the Link to Search button in the search toolbar to display the URL that creates the search. Users who execute this search via URL may still be required to authenticate through the Tealeaf Portal. See "Searching Session Data" in the *IBM Tealeaf cxImpact User Manual*.

Overview

All parameters must be URL-encoded.

Required Arguments: template

Argument
Value

`template`
The ID of the template to be loaded. The user must have permission to view this template. The IDs for templates can be found on the search template config page with the **Edit** button for templates.

`fieldX` Either the active search keyword or the archive search keyword (e.g. request) for a search field. The first matching search field will be selected from the template and used. Knowledge of acceptable keywords is required, so we have to provide a master list somewhere.

`opX` The operator value for the field

valueAX
The first value to populate the field with. Text inputs only use valueA - others such as form fields, or numerical fields with the `between` operator can use valueB

valueBX
The second value to populate the field with

For `field`, `op`, `valueA` and `valueB`, the `X` is a numerical index starting at `0`. For example, `field0=request`, `op0=1`, `valueA0=tealeaf` will find a search field in the template that uses the `request` keyword, set the operator to `includes`, and populate it with the text `tealeaf`.

Required Arguments: template

Argument
Value

datespan
The numeric constant representing the dates to search on.

startdate
If datespan is set to -1, this is the ISO-format date for the beginning date of the range (e.g. 2007-01-02 is Jan 02, 2007)

enddate
If datespan is set to -1, this is the ISO-format date for the end date of the range (e.g. 2007-01-02 is Jan 02, 2007)

starttime
If datespan is set to -1, this is the 6-digit, 24-hour time for the beginning of the range, in HHMMss format (e.g. 003015 is 00:30:15 or 12:30:15 AM)

endtime
If datespan is set to -1, this is the 6-digit, 24-hour time for the end of the range, in HHMMss format (e.g. 233059 is 23:30:59 or 11:30:59 PM)

scope The scope of the search.

sort How the results will be sorted

timeout
The maximum allotted time for an archive search, in seconds. The search will be stopped once this amount of time has elapsed, but more time may be required to collate the sessions found.

autostop
Specifies the limit at which an archive search will be stopped. When X matching sessions are found, the search is stopped.

dosearch
If set to 1, the search will be performed once all the fields have been processed and populated. This mode takes you to the search builder, so the search results will appear as a second breadcrumb. This allows a search to be performed, but the user will still be able to go back and tweak the constraints to refine the search.

Operator Values

The following values may be included in the `opX` parameter:

Value **Description**

1 Includes/contains (active/archive)

-1 NOT includes/contains

2 Equal (numerical)

-2 NOT equal

3 Between (numerical)

-3 NOT between

4 Greater than

5 Less than

CX User Administration

This section details user administration for IBM Tealeaf CX.

Note: If you have licensed IBM Tealeaf cxResults, you may choose to configure separate user and group permissions for IBM Tealeaf cxResults users, which are defined through IBM Tealeaf CX User Administration. See "Managing Users and Groups for cxResults" in the *IBM Tealeaf cxResults Administration Manual*.

To access IBM Tealeaf CX User Administration:

1. Login to the Portal as an administrator.
2. From the Portal menu, select **Tealeaf** > **Portal Management**.
3. The Portal Management page is displayed. In the left panel, click **IBM Tealeaf CX User Administration**.
4. Select the category that applies:
   - "Users"
   - "Groups" on page 169
   - "Current Users" on page 174

Note: Changes to settings saved in the Portal Management page are not applied to the Portal application until the page is refreshed through your browser.

Users

To see a list of all IBM Tealeaf CX users, click **Users**. To make changes to a user, select it and click one of the buttons above the list to make changes.

New / Edit

This section explains how to add and edit details for new users.

To create a new user, click **New**. To edit an existing one, select the user, and click **Edit**.

Some user properties can be set to **Primary Group Default**, which sets the user property to the value defined for the group specified as the user's Primary User Group.

Setting Description

Active If this box is deselected, the user cannot login or receive emails.

ID This system-generated number is used internally to identify the account.

Username
The login name for the user.

Note: Tealeaf recommends using only alphanumeric characters when specifying names used by the Portal application. Some non-alphanumeric characters may be blocked or may cause issues for some Tealeaf components.

NT Username
If NT authentication is in use, this value is the NT/AD user name for the user.

NT Domain
If NT authentication is in use, this value is the NT/AD domain for the user.

Email Enter a valid email address for delivery of results set analysis, scheduled dashboards, alerts, and other Tealeaf content.

Time Zone (used in Search)
Hourly report values are displayed in the user's timezone, but the daily rollup of reporting data is controlled by the system timezone. It is best to set this value to the same as the system timezone to avoid confusion.

- This time zone setting applies only to the search activities. For individual searches, the time zone can be configured in the Options screen. See "Searching Session Data" in the *IBM Tealeaf cxImpact User Manual*.

System Locale
Select the preferred system locale setting to configure the Tealeaf Portal to use the regional locale settings for the user. The system locale setting defines how currency, numeric values, date, and time are formatted are displayed.

This setting defaults to the **Primary Group Default** at the user level. The administrator can manage this setting by selecting a system locale for each user group.

Date Format
Select the format, such as `mm/dd/yyyy`, from the list. The **Date Format** setting defaults to the **Primary Group Default** date. The **Primary Group Default** uses the selected date format setting that is assigned to the user group. If you want to specify a custom system locale for a user, you can configure the **System Locale** setting then, select **Use system locale** for the date format. If you want to use a date format that is unrelated to the system locale setting, you can select a custom date format from the list.

Language
The language the Portal should display for the user.

Password Expires
The number of days after which the Portal does not permit a user from continuing after login until the user configured a new password.

Start Page
The list of available start pages depends on the user's group and menu permissions.

Portal Navigation Menu
Enables or disables the user's ability to see the top navigation menu in the Portal.

Log User Out If Idle
When enabled, the user is logged out if the session is idle for a period longer than the session timeout length.

- You can change the timeout length through Report configuration. See "Configuring the Report Server" in the *IBM Tealeaf CX Configuration Manual*.

Allow Report Template Import
When enabled, the user is permitted to import report templates. When set to `Primary Group Default`, the user inherits the default setting for the primary group.

- All users can export report templates regardless of this setting.

Create Session Segment for Completed Session Search
By default, every search the user performs on completed sessions creates a session segment that remains available for 24 hours.

Create Combined Session Segment
Enables or disables the Build Combined Segment feature in the Manage Session Segments page for the IBM Tealeaf cxReveal user. See "Managing Session Segments" in the *IBM Tealeaf cxResults User Manual*.

Lock Replay Mode
Enables or disables the user's ability to change his or her Replay Mode.

Replay Mode
The following values are available:

- `RTV` - Launch the IBM Tealeaf CX RealiTea Viewer desktop application to replay sessions. The app must be already installed on the user's desktop.
- `Browser` - Replay sessions inside the browser.
- `Prompt on Replay` - Open a dialog in which the user can choose RTV or Browser for every replay.

Dashboard Detail Mode
Display drill-down information in the same browser window or display it in a new one.

Default Active Search Template
The default search template to load when the user chooses to search for active sessions.

Default Completed Search Template
The default search template to load when the user chooses to search for completed sessions.

Default All Sessions Search Template
The default search template to load when the user chooses to search for all sessions.

Search Box for Online Help
Enable or disable the Online Help search box available at the top of each Portal page.

TMS Access Level
Determine the TMS Access Level for the user. You can choose from full access or restricted access.

Open BBR In A New Window
When enabled, Browser Based Replay sessions are displayed in a new browser window.

Primary User Group
The primary group to use for group level settings.

Member Of
Lists all groups of which the user is currently a member.

To apply changes, click **Save**.

Password

To change a user's password, select a user and click **Password**. Enter the new user's password twice. To apply the change, click **Save**.

Force Password Change

To force the user to change his or her password, select a user and click **Force Password**. When the user next logs in, the user is required to change the account password before using the Portal.

Clear Password History

Reset the user's password history, allowing the user to re-use passwords that might have conflicted with the Portal's password usage policy.

Assign Groups

To assign the user to one or more groups, select a user and click **Assign Groups**.

To switch a user's role, remove the user from all groups, and then add a group for the user's new role.

Manage NT Users

When NT authentication is enabled in the Portal, Tealeaf administrators can use this feature to create and disable Portal accounts based on Windows domain logins detected by the Portal.

- For more information on authentication, see “Authentication” on page 329.
- NT authentication is enabled through Search Server configuration. See "Configuring the Search Server" in the *IBM Tealeaf CX Configuration Manual*.
- If desired, you can configure the Portal to automatically login and create accounts for any NT user who visits the Tealeaf Portal. See “CX Settings” on page 192.

Left pane: NT users who do not have Portal accounts **Right pane:** Current Portal accounts

- To mark an NT user for Portal account creation, select the user, and click the >> arrow.
- To disable a Portal account, select the user, and click the << arrow.
- To save these changes, click **Save**. The marked accounts are created or disabled.

View Activity

For the selected user, you can review the available activity reports in the User Activity page in the Portal.

- See “Monitoring User Activity” on page 57.

Delete

Deletes the selected user's profile completely, including any dashboards and reports created by that user.

Note:

1. When you delete a user account, any scorecard schedules or dashboards created by the user are removed. To avoid this deletion, you may change the password on the user account instead.
2. The Admin user cannot be deleted.

Groups

You can create new user groups or edit permissions for existing groups.

To see a list of users in the selected group, select a group in the User Administration screen and click **Users**.

New / Edit

This section explains how to add and edit groups.

To create a new group, click **New**. To edit an existing group, select the group and click **Edit**.

User groups have the following properties. All settings except for the Group Name setting can be inherited by individual users who have the selected group defined as their Primary User Group.

Setting Description

ID This system-generated number is used internally to identify the group.

Group Name
The name of the group.

Note: Tealeaf recommends using only alphanumeric characters when specifying names used by the Portal application. Some non-alphanumeric characters may be blocked or may cause issues for some Tealeaf components.

Default Time Zone (used in Search)
Hourly report values are displayed in the user's timezone, but the daily rollup of reporting data is controlled by the system timezone. It is best to set this value to the same as the system timezone to avoid confusion.

- This time zone setting applies primarily to the affected users' search activities. Other Portal activities utilize the Tealeaf system time zone. See "Configuring the System Timezone" in the *IBM Tealeaf CX Configuration Manual*.

Default System Locale
Select your preferred system locale. This setting applies the selected regional formatting to dates, time, and numeric values that are displayed to the user within the Tealeaf Portal.

Default Date Format
Select the format, such as `mm/dd/yyyy`, from the list. By default, the date format is set to the date format of the configured system locale.

Default Language
The language the Portal should display for the user group.

Default Start Page
The list of available start pages depends on the user's group and menu permissions.

Default Password Expiration
This option requires all users to change their passwords every N days. After the password expires, the portal does not let the user continue after login until a new password has been set.

Password Expiration Warning Period
The number of days before password expiration for which the system warns users that their passwords are about to expire.

Default Portal Navigation Menu
Enable or disable the group's ability to see the top navigation menu in the Portal.

Default Log User Out If Idle
When enabled, members of the group are automatically logged out if their sessions are idle for a period longer than the session timeout length.

- You can change the timeout length through Report configuration. See "Configuring the Report Server" in the *IBM Tealeaf CX Configuration Manual*.

Default Allow Report Template Import
When enabled, this setting allows all users of the group to import report templates by default. This setting can be overridden by settings to individual user accounts.

- By default, admin accounts are configured to allow report import.
- All users can export report templates regardless of this setting.

Default Create Session Segment for Completed Session Search
Enable or disable the default length of time for which session segments remain available.

Default Create Combined Session Segment
Enable or disable the IBM Tealeaf cxReveal group's ability to use the Build Combined Segment feature in the Manage Session Segments page. See "Managing Session Segments" in the *IBM Tealeaf cxResults User Manual*.

Default Lock Replay Mode
Enable or disable the group's ability to change Replay Mode.

Default Replay Mode
The following options are available:

- `RTV` - Launch the IBM Tealeaf CX RealiTea Viewer desktop application to replay sessions. The app must be installed on the local desktops of users in the group.
- `Browser` - Replay sessions inside the browser.
- `Prompt on Replay` - Pop up a dialog to choose RTV or Browser for every replay.

Default Dashboard Detail Mode
When users drill-down into a dashboard, this option defines how the details are displayed.

Default Active Search Template
The default search template to load when the members of the group choose to search for active sessions.

Default Completed Search Template
The default search template to load when the members of the group choose to search for completed sessions.

Default All Sessions Search Template
The default search template to load when the members of the group choose to search for all sessions.

Search Box for Online Help
Enable or disable the Online Help search box available at the top of each Portal page for the group.

TMS Access Level
Determine the TMS Access Level for the group. You can choose from full access or restricted access.

Open BBR In a New Window
When enabled, sessions in Browser Based Replay are displayed in a new window for the group.

Assign Users

To add users to the selected group, click **Assign Users**. The list is populated with users who have the proper role (IBM Tealeaf CX, IBM Tealeaf cxView, or IBM Tealeaf cxReveal) yet do not have a group assignment.

To re-assign a user to a new group, remove all group assignments from the user and then assign the user to the new group to match the user's new role.

Menu Profile

The Menu Profile for a user group determines the available menu items that members of the group can access. It also controls the reports that can be viewed on member Dashboards. For example, if a group's menu profile does not include access to Analyze > Report Builder, its members cannot view the Report Builder page.

- Users of some Tealeaf products, such as IBM Tealeaf cxView and IBM Tealeaf cxReveal, have access to a subset of Portal menu options and cannot access other menu options due to their Tealeaf license.

  Note: After you modify a user group's menu profile, you should verify that the users in the group can still access their default page. This page is defined at the group level in the `Default Start Page` setting and is also available for individual user configuration. Users who cannot access their default page are forced to review their settings. See "My Settings" in the *IBM Tealeaf cxImpact User Manual*.

Administration Profile

For groups that are granted access to the Portal Management page, the Administration Profile controls the portions of the page to which group members have access. The items in the list correspond to the sections on the Administration page sidebar.

Search Profile

For groups that are granted access to the Session Search page, the Search Profile controls both the search options and search templates to which they have access. If a user has no available search options, the Options section is hidden entirely on the Session Search page.

- See "Searching Session Data" in the *IBM Tealeaf cxImpact User Manual*.
- **IBM Tealeaf cxResults:** For more information on searching for visitors, see "Searching for Visitors" in the *IBM Tealeaf cxResults User Manual*.
- For more information on configuring search templates, see "Configuring Search Templates" in the *IBM Tealeaf cxImpact Administration Manual*.

Setting Description

Date Span Selector
When selected, users in this group can access the Date Span selector, which enables the selection of specific date ranges.

Servers Selector
When selected, users in this group can access the Servers selector, which enables the selection of specific servers.

Template Selector
When selected, users in this group can access the Template selector drop-down, which enables the selection of non-default search templates.

Session List Template
When selected, users in this group can access the Session List Template selector, which enables the selection of non-default session list templates to display search results.

Search Scope
When selected, users in this group can access the Search Scope selector, which enables the selection of the page-level and session-level scoping options.

Search Results Sort
When enabled, users in this group can choose to sort the search query to return results prioritized by a field value selected in the Sort Results By drop-down in Search Options.

- See "Searching Session Data" in the *IBM Tealeaf cxImpact User Manual*.

Search Timeout
When selected, the default search timeout limit is imposed on searches executed by this group.

- Searches that require more time than this limit are stopped before completion. This setting can be configured in TMS. See "Configuring the Search Server" in the *IBM Tealeaf CX Configuration Manual*.

  Note: Avoid deselecting this option. Allowing unlimited searches can cause Search Server to crash.

Search Autostop Limit
When selected, the default autostop limit is imposed on searches executed by this group.

Saved Searches
When selected, users in this group can open and execute searches that have been saved.

Browser Replay Profile

For groups that are granted access to Browser Based Replay (BBR), the Browser Replay Profile controls the actions and permissions to which group members have access while using BBR. The replay and navigation sections of BBR are available to all users, but other functions are controlled by the Browser Replay Profile.

- See "Browser Based Replay Interface" in the *IBM Tealeaf cxImpact User Manual*
- See "BBR Options" in the *IBM Tealeaf cxImpact User Manual*
- You can configure rules for BBR replay through the Portal Management page. See "BBR Replay Rules" in the *IBM Tealeaf cxImpact User Manual*.

Table 8. Browser replay profile

Setting	Description
`View Request`	When selected, users in this group can access BBR Request view.
`View Response`	When selected, users in this group can access BBR Response view.
`View Events`	When selected, users in this group can access the Events pane.
`View URL Fields`	When selected, users in this group can view the URL fields in the left panes.
`View Gen Time`	When selected, the Gen Time statistic is displayed in the page statistics bar. • See "Browser Based Replay Interface" in the *IBM Tealeaf cxImpact User Manual*.
`View Net Trip Time`	When selected, the Net Trip Time statistic is displayed in the page statistics bar. • See "Browser Based Replay Interface" in the *IBM Tealeaf cxImpact User Manual*.
`View Round Trip Time`	When selected, the Round Trip Time statistic is displayed in the page statistics bar. • See "Browser Based Replay Interface" in the *IBM Tealeaf cxImpact User Manual*.
`View Bytes`	When selected, the size of the response in bytes is displayed in the page statistics bar. • See "Browser Based Replay Interface" in the *IBM Tealeaf cxImpact User Manual*.
`View Status Code`	When selected, the status code of the response is displayed in the page statistics bar. • See "Browser Based Replay Interface" in the *IBM Tealeaf cxImpact User Manual*.

Table 9. Actions

Setting	Description
`Enable Searching`	When selected, users in this group can copy text selections and use them to execute searches of completed sessions from within BBR.
`Enable Email Sessions`	When selected, users in this group can email links to sessions opened in BBR.

Table 10. Settings

Setting	Description
Merge Session Fragments	When selected, this option is available in the drop-down menu. When it is selected, users in this group can configure and execute the merging of session fragments. • Merging session fragments enables the stitching together of fragments of an individual session that was splintered by timeout, traffic issues, or other factors. The configured behavior can be overridden through the interface. See "BBR Options" in the *IBM Tealeaf cxImpact User Manual*.
Report Javascript Errors	When selected, this option is available in the drop-down menu. When it is selected, JavaScript errors are reported to the BBR user.
Resize Replay Pane To Match Sessions	When selected, this option is available in the drop-down menu. When it is selected, the replay pane in BBR is automatically resized to match the screen dimensions that are used by the visitor's user agent.
Advanced Options	Select this option to enable the ability to load Tealeaf archive files into BBR for replay. **Note:** TLA files loaded through BBR do not have On-Demand Privacy rules applied to them and may pose a security risk if this feature is enabled. See "On-Demand Privacy" in the *IBM Tealeaf CX Configuration Manual* • Replay rules are applied to loaded TLA files. • For more information about these settings in BBR, see "BBR Options" in the *IBM Tealeaf cxImpact User Manual*.
Select Page Name	This drop-down menu allows you to determine if you want to Show URL, Show Page Title, or Show TLT_URL. You can set the default page name here across all sessions.

Delete

Deletes the selected group. Members of the group are not deleted.

Current Users

Displays a list of the currently active users of the specified role.

- To review a user's recent activity, click the username. The Audit page is displayed.

cxReveal User Administration

The IBM Tealeaf cxImpact, IBM Tealeaf cxReveal and IBM Tealeaf cxView user administration sections are very similar, but there are differences in the settings for each. The three types of users and groups they manage have very different roles.

To access IBM Tealeaf cxReveal User Administration:

1. Login to the Portal as an administrator.
2. From the Portal menu, select **Tealeaf** > **Portal Management**.
3. The Portal Management page is displayed. In the left panel, click **IBM Tealeaf cxReveal User Administration**.
4. Select the category that applies:
   a. “Users” on page 175
   b. “Groups” on page 178

c. “Current Users” on page 183

Note: Changes to settings saved in the Portal Management page are not applied to the Portal application until the page is refreshed through your browser.

Note: If you have installed this product as part of a new installation of Tealeaf, access to product features through the Portal are provided to the `ADMIN` group automatically. Automatic access is also granted if you have licensed this product for an existing Tealeaf solution and have not modified your `ADMIN` group permissions since original installation.

- If you have modified the `ADMIN` group and are deploying this product to an existing Tealeaf solution, you must update menu permissions manually for that group.
- For all other users and groups, you must configure the appropriate permissions.

Note: By default, IBM Tealeaf cxReveal accounts created in a new install are assigned to the `cxReveal User` group, since most IBM Tealeaf cxReveal accounts require only this level of access. After the account is created, admin-level permissions can be assigned through this screen to individual accounts. Account assignments and permissions are not altered on upgrade.

Users

Click **Users** to see a list of all IBM Tealeaf cxReveal users.

Select a user and click one of the buttons above the list to make changes.

New / Edit

Click **New** to create a new user, or select an existing user from the list and click **Edit** to open the user properties dialog

Some of the properties can be set to **Primary Group Default**, which means the user inherits the setting from the user's primary group, which is specified in the Primary User Group property.

Setting Description

`Active` If this box is deselected, the user cannot login or receive emails.

`ID` This system-generated number is used internally to identify the account.

`Username`
The login name for the user.

Note: Tealeaf recommends using only alphanumeric characters when specifying names used by the Portal application. Some non-alphanumeric characters may be blocked or may cause issues for some Tealeaf components.

`NT Username`
If NT authentication is in use, this value is the NT/AD user name for the user.

`NT Domain`
If NT authentication is in use, this value is the NT/AD domain for the user.

`Email` Enter a valid email address for delivery of results set analysis, scheduled dashboards, alerts, and other Tealeaf content.

Time Zone (used in Search)
Hourly report values are displayed in the user's timezone, but the daily rollup of reporting data is controlled by the system timezone. It is best to set this value to the same as the system timezone to avoid confusion.

- This time zone setting applies primarily to the affected users' search activities. Other Portal activities utilize the Tealeaf system time zone. See "Configuring the System Timezone" in the *IBM Tealeaf CX Configuration Manual*.

System Locale
Select the user's preferred system locale. The system locale setting defines how currency, numeric values, date, and time are formatted are displayed.

This setting defaults to the **Primary Group Default** at the user level. The administrator can manage this setting by selecting a system locale for each user group.

Date Format
Select the format, such as `mm/dd/yyyy`, from the list. The **Date Format** setting defaults to the **Primary Group Default** date. The **Primary Group Default** uses the selected date format setting that is assigned to the user group. If you want to specify a custom system locale for a user, you can configure the **System Locale** setting then, select **Use system locale** for the date format. If you want to use a date format that is unrelated to the system locale setting, you can select a custom date format from the list.

Language
The language the Portal should display for the user.

Password Expires
The number of days after which the Portal does not permit a user from continuing after login until the user has configured a new password.

Start Page
The list of available start pages depends on the user's group and menu permissions.

Portal Navigation Menu
Enables or disables the user's ability to see the top navigation menu in the Portal.

Log User Out If Idle
When enabled, the user is logged out if the session is idle for a period longer than the session timeout length.

- You can change the timeout length through Report configuration. See "Configuring the Report Server" in the *IBM Tealeaf CX Configuration Manual*.

Allow Report Template Import
When enabled, the user is permitted to import report templates. When set to `Primary Group Default`, the user inherits the default setting for the primary group.

- All users can export report templates regardless of this setting.

Create Combined Session Segment
Enables or disables the Build Combined Segment feature in the Manage Session Segments page for IBM Tealeaf cxReveal users. See "Managing Session Segments" in the *IBM Tealeaf cxResults User Manual*.

Lock Replay Mode
Enables or disables the user's ability to change his or her Replay Mode.

Replay Mode
The following values are available:

- `RTV` - Launch the IBM Tealeaf CX RealiTea Viewer desktop application to replay sessions. The app must be already installed on the user's desktop.
- `Browser` - Replay sessions inside the browser.
- `Prompt on Replay` - Open a dialog in which the user can choose RTV or Browser for every replay.

Default Active Search Template
The default search template to load when the user chooses to search for active sessions.

Default Completed Search Template
The default search template to load when the user chooses to search for completed sessions.

Default All Sessions Search Template
The default search template to load when the user chooses to search for all sessions.

Search Box for Online Help
Enable or disable the Online Help search box available at the top of each Portal page.

Open BBR In A New Window
When enabled, sessions loaded in Browser Based Replay are displayed in a new browser window.

Primary User Group
The primary group to use for group level settings.

Member Of
Lists all groups of which the user is currently a member.

Click **Save** to apply your changes, or click **Cancel** to quit without applying them.

Password

To change a password, select a user and click **Password**. Enter the new user's password twice in the **New Password** and **Confirm New Password** fields. Click **Save** to apply the change.

Force Password Change

To force the user to change his or her password, select a user and click **Force Password Change**. The next time the user logs in, he or she is required to change the account password before using the Portal.

Clear Password History

Reset the user's password history, allowing the user to re-use passwords that might have conflicted with the Portal's password usage policy.

Assign Groups

Select a user and click **Assign Groups** to assign the user to one or more groups. To switch a user's role, remove the user from all groups, and then add a group for the user's new role.

Manage NT Users

When NT authentication is enabled in the Portal, Tealeaf administrators can use this feature to create and disable Portal accounts based on Windows domain logins detected by the Portal.

- For more information on authentication, see "Authentication" in the *IBM Tealeaf cxImpact Administration Manual*.
- NT authentication is enabled through Search Server configuration. See "Configuring the Search Server" in the *IBM Tealeaf CX Configuration Manual*.
- If desired, you can configure the Portal to automatically login and create accounts for any NT user who visits the Tealeaf Portal. See "CX Settings" in the *IBM Tealeaf cxImpact Administration Manual*.

Left pane: NT users who do not have Portal accounts
Right pane: Current Portal accounts

- To mark an NT user for Portal account creation, select the user, and click the >> arrow.
- To disable a Portal account, select the user, and click the << arrow.
- To save these changes, click **Save**. The marked accounts are created or disabled.

View Activity

This button links to the Audit page, which lists Portal activity for the selected user. For more information see the "Audit page" section.

Delete

Deletes the selected user's profile completely, as well as any dashboards and scorecards created by that user.

Note:

1. When you delete a user account, any scorecard schedules or dashboards created by the user are removed. To avoid this deletion, you may change the password on the user account instead.
2. The Admin user cannot be deleted.

Groups

Under IBM Tealeaf cxReveal User Administration, click **Groups** to display a list of the groups.

New / Edit

Click **New** to create a new group, or select a group from the list and click **Edit** to open the group properties dialog.

User groups have the following properties. All settings except for the Group Name setting can be inherited by users who have the selected group set as their Primary User Group.

Setting Description

ID This system-generated number is used internally to identify the group.

Group Name
The name of the group.

Note: Tealeaf recommends using only alphanumeric characters when specifying names used by the Portal application. Some non-alphanumeric characters may be blocked or may cause issues for some Tealeaf components.

Default Time Zone (used in Search)
Hourly report values are displayed in the user's timezone, but the daily

rollup of reporting data is controlled by the system timezone. It is best to set this value to the same as the system timezone to avoid confusion.

- This time zone setting applies primarily to the affected users' search activities. Other Portal activities utilize the Tealeaf system time zone. See "Configuring the System Timezone" in the *IBM Tealeaf CX Configuration Manual*.

Default System Locale
Select your preferred system locale. This setting applies the selected regional formatting to dates, time, and numeric values that are displayed to the user within the Tealeaf Portal.

Default Date Format
Select the format, such as `mm/dd/yyyy`, from the list. By default, the date format is set to the date format of the configured system locale.

Default Language
The language the Portal should display for the user group.

Default Start Page
The list of available start pages depends on the user's group and menu permissions.

Default Password Expiration
This option requires all users to change their passwords every N days. After the password expires, the portal does not let the user continue after login until a new password has been set.

Password Expiration Warning Period
The number of days before password expiration for which the system warns users that their passwords are about to expire.

Default Portal Navigation Menu
Enable or disable the group's ability to see the top navigation menu in the Portal.

Default Log User Out If Idle
When enabled, members of the group are automatically logged out if their sessions are idle for a period longer than the session timeout length.

- You can change the timeout length through Report configuration. See "Configuring the Report Server" in the *IBM Tealeaf CX Configuration Manual*.

Default Allow Report Template Import
When enabled, this setting allows all users of the group to import report templates by default. This setting can be overridden by settings to individual user accounts.

- By default, admin accounts are configured to allow report import.
- All users can export report templates regardless of this setting.

Default Create Combined Session Segment
Enable or disable the IBM Tealeaf cxReveal group's ability to use the Build Combined Segment feature in the Manage Session Segments page. See "Managing Session Segments" in the *IBM Tealeaf cxResults User Manual*.

Default Lock Replay Mode
Enable or disable the group's ability to change Replay Mode.

Default Replay Mode
The following options are available:

- `RTV` - Launch the IBM Tealeaf CX RealiTea Viewer desktop application to replay sessions. The app must be installed on the local desktops of users in the group.
- `Browser` - Replay sessions inside the browser.
- `Prompt on Replay` - Pop up a dialog to choose RTV or Browser for every replay.

`Default Active Search Template`
The default search template to load when the members of the group choose to search for active sessions.

`Default Completed Search Template`
The default search template to load when the members of the group choose to search for completed sessions.

`Default All Sessions Search Template`
The default search template to load when the members of the group choose to search for all sessions.

`Search Box for Online Help`
Enable or disable the Online Help search box available at the top of each Portal page for the group.

`Open BBR In New Window`
When enabled, sessions loaded in Browser Based Replay are displayed in a new window for the group.

Assign Users

Click **Assign Users** to add users to the selected group.

The list is populated with users who are of the correct role (IBM Tealeaf cxImpact, IBM Tealeaf cxView, or IBM Tealeaf cxReveal) and users who are members of absolutely no other groups.

Including users without groups allows you to transition users between roles. To move a user to a new role, remove all group assignments from the user, and then assign the user to another group that matches the new role.

Menu Profile

The Menu Profile for a user group determines the menu options that a group's members can access. It also controls which reports can be viewed on the dashboards of a group's members. For example, if a group does not have permission in the menu profile for Analyze > Report Builder, its members cannot view the Report Builder page or any reports generated by other Tealeaf users.

Note: IBM Tealeaf cxReveal provides a set of menu options that has been reduced from the menu options available to IBM Tealeaf cxImpact users. For more information on available menu options in IBM Tealeaf cxImpact, see "Portal Navigation" in the *IBM Tealeaf cxImpact User Manual*.

Note: After you modify a user group's menu profile, you should verify that the users in the group can still access their default page. This page is defined at the group level in the `Default Start Page` setting and is also available for individual user configuration. Users who cannot access their default page are forced to review their settings. See "My Settings" in the *IBM Tealeaf cxImpact User Manual*.

Administration Profile

For groups granted access to the Portal Management page, the Administration Profile controls the portions of the page that they can see. The items in the list correspond to the sections on the Administration page sidebar.

Search Profile

For groups granted access to the Session Search page, the Search Profile controls the search options members that they can see and the search templates that they are allowed to use. If a user has no search options available, the options section is hidden entirely on the Session Search page.

Browser Replay Profile A

For groups that are granted access to Browser Based Replay (BBR), the Browser Replay Profile controls the actions and permissions to which group members have access while using BBR. The replay and navigation sections of BBR are available to all users, but other functions are controlled by the Browser Replay Profile.

- See "Browser Based Replay Interface" in the *IBM Tealeaf cxImpact User Manual*
- See "BBR Options" in the *IBM Tealeaf cxImpact User Manual*
- You can configure rules for BBR replay through the Portal Management page. See "BBR Replay Rules" in the *IBM Tealeaf cxImpact User Manual*.

Browser replay profile B:

Setting Description

View Request
When selected, users in this group can access BBR Request view.

View Response
When selected, users in this group can access BBR Response view.

View Events
When selected, users in this group can access the Events pane.

View URL Fields
When selected, users in this group can view the URL fields in the left-hand panes.

View Gen Time
When selected, the Gen Time statistic is displayed in the page statistics bar.

See "Browser Based Replay Interface" in the *IBM Tealeaf cxImpact User Manual*.

View Net Trip Time
When selected, the Net Trip Time statistic is displayed in the page statistics bar.

See "Browser Based Replay Interface" in the *IBM Tealeaf cxImpact User Manual*.

View Round Trip Time
When selected, the Round Trip Time statistic is displayed in the page statistics bar.

See "Browser Based Replay Interface" in the *IBM Tealeaf cxImpact User Manual*.

View Bytes
When selected, the size of the response in bytes is displayed in the page statistics bar.

See "Browser Based Replay Interface" in the *IBM Tealeaf cxImpact User Manual*.

View Status Code
When selected, the status code of the response is displayed in the page statistics bar.

See "Browser Based Replay Interface" in the *IBM Tealeaf cxImpact User Manual*.

Actions:

Setting Description

Enable Searching
When selected, users in this group can copy text selections and use them to execute searches of completed sessions from within BBR.

Enable Email Sessions
When selected, users in this group can email links to sessions opened in BBR.

Settings:

Setting Description

Use Page Titles Instead of URLs
When selected, this option is available in the drop-down menu. When it is selected, the Page Title of each page is listed in the Navigation panel, instead of the URL.

Merge Session Fragments
When selected, this option is available in the drop-down menu. When it is selected, users in this group can configure and execute the merging of session fragments.

Merging session fragments enables the stitching together of fragments of an individual session that has been splintered by timeout, traffic issues, or other factors. The configured behavior can be overridden through the interface. See "BBR Options" in the *IBM Tealeaf cxImpact User Manual*.

Report Javascript Errors
When selected, this option is available in the drop-down menu. When it is selected, Javascript errors are reported to the BBR user.

Resize Replay Pane To Match Sessions
When selected, this option is available in the drop-down menu. When it is selected, the replay pane in BBR is automatically resized to match the screen dimensions used by the visitor's user agent.

Advanced Options
Select this option to enable the ability to load Tealeaf archive files into BBR for replay.

Note: TLA files loaded through BBR do not have On-Demand Privacy rules applied to them and may pose a security risk if this feature is enabled. See "On-Demand Privacy" in the *IBM Tealeaf CX Configuration Manual*.

- Replay rules are applied to loaded TLA files.
- For more information on these settings in BBR, see "BBR Options" in the *IBM Tealeaf cxImpact User Manual*.

Delete

Deletes the selected group. The group's members are not deleted.

Current Users

Displays a list of the currently active users of the specified role. The user's name is a link to the Audit page, which displays the user's activity.

Portal Menu Bar

In many implementations, IBM Tealeaf cxReveal users do not require access to a Portal menu bar, as users' primary functions are searching for sessions and then replaying found sessions. Depending on your user account permissions, you may not be able to see a menu inside the Portal.

If needed, Tealeaf administrators can configure the menu profile for individual IBM Tealeaf cxReveal users or IBM Tealeaf cxReveal user groups to include a menu bar in the Portal.

Note: A user's menu profile is defined as the union of the menu profiles of all user groups to which the user belongs. For example, if the user belongs to two user groups, the menu profiles of each group are applied to the user.

- If a specific permission is enabled for one group and disabled for another group to which the user belongs, then the menu item is enabled.

For cxReveal Groups

To make menu bar changes for IBM Tealeaf cxReveal groups, please complete the following steps.

1. Login to the Portal as an administrator.
2. From the menu, select **Tealeaf** > **Portal Management**.
3. In the left panel, click the IBM Tealeaf cxReveal User Administration link.
4. Click the Groups link.
5. Select the group whose menu profile you'd like to change.
6. From the Default Portal Navigation Menu setting, you may select one of the following options:
   - `Enable` - Enable the Portal menu for this user group.
   - `Disable` - Disable the Portal menu for this user group.
     - IBM Tealeaf cxReveal users in this group are placed in the page specified in the Default Start Page setting. Typically, this value is set to `Default Search Template` for IBM Tealeaf cxReveal users.
7. To configure a menu profile for the group, click **Menu Profile**.
8. In the Menu navigation tree, you can select the menus and menu items to display to members of this group.
9. To save changes, click **Save**.
   - See "Menu Profile" on page 180.

For Individual cxReveal Users

To make menu bar changes for individual IBM Tealeaf cxReveal users, please complete the following steps.

1. Login to the Portal as an administrator.
2. From the menu, select **Tealeaf** > **Portal Management**.
3. In the left panel, click the IBM Tealeaf cxReveal User Administration link.

4. Click the Users link.
5. Select the user whose menu profile you'd like to enable.
6. Click **Edit**.
7. From the Portal Navigation Menu setting, you may select one of the following options:
   - `Primary Group Default` - This user inherits its menu profile settings from the primary user group configured for the user. This group is listed under the Primary User Group setting for the user account.
   - `Enable` - Enables a menu profile for this individual user.
   - `Disable` - Disables a menu profile for this individual user.
8. To save changes, click **Save**.
   - See "Users" on page 175.

cxView User Administration

Users and groups are managed in the Portal through the User Administration sections. As there are different roles for users, each section applies only to the users of that type, although they all share properties.

Note: If you have installed this product as part of a new installation of Tealeaf, access to product features through the Portal are provided to the `ADMIN` group automatically. Automatic access is also granted if you have licensed this product for an existing Tealeaf solution and have not modified your `ADMIN` group permissions since original installation.

- If you have modified the `ADMIN` group and are deploying this product to an existing Tealeaf solution, you must update menu permissions manually for that group.
- For all other users and groups, you must configure the appropriate permissions.

To accessIBM TealeafcxView User Administration:

1. Login to the Portal as an administrator.
2. From the Portal menu, select **Tealeaf** > **Portal Management**.
3. The Portal Management page is displayed. In the left panel, click **IBM Tealeaf cxView User Administration**.
4. Select the category that applies:
   - "Users"
   - "Groups" on page 187
   - "Current Users" on page 189

Note: Changes to settings saved in the Portal Management page are not applied to the Portal application until the page is refreshed through your browser.

Users

Under IBM Tealeaf cxView User Administration, click **Users** to display a list of the IBM Tealeaf cxView users.

New / Edit

Click **New** to create a new user, or select an existing user from the list and click **Edit** to open the user properties dialog.

Some of the properties can be set to **Primary Group Default**, which means the user inherits the setting from the group specified as the **Primary User Group** property.

Setting Description

Active If this box is deselected, the user cannot login or receive emails.

ID This system-generated number is used internally to identify the account.

Username
The login name for the user.

Note: Tealeaf recommends using only alphanumeric characters when specifying names used by the Portal application. Some non-alphanumeric characters may be blocked or may cause issues for some Tealeaf components.

NT Username
If NT authentication is in use, this value is the NT/AD user name for the user.

NT Domain
If NT authentication is in use, this value is the NT/AD domain for the user.

Email Enter a valid email address for delivery of results set analysis, scheduled dashboards, alerts, and other Tealeaf content.

Time Zone (used in Search)
Hourly report values are displayed in the user's timezone, but the daily rollup of reporting data is controlled by the system timezone. It is best to set this value to the same as the system timezone to avoid confusion.

- This time zone setting applies primarily to the affected users' search activities. Other Portal activities utilize the Tealeaf system time zone. See "Configuring the System Timezone" in the *IBM Tealeaf CX Configuration Manual*.

Date Format
Select the format, such as `mm/dd/yyyy`, from the drop-down list.

Language
The language the Portal should display for the user.

Password Expires
The number of days after which the Portal does not permit a user from continuing after login until the user has configured a new password.

Start Page
The list of available start pages depends on the user's group and menu permissions.

Portal Navigation Menu
Enables or disables the user's ability to see the top navigation menu in the Portal.

Log User Out If Idle
When enabled, the user is logged out if the session is idle for a period longer than the session timeout length.

- You can change the timeout length through Report configuration. See "Configuring the Report Server" in the *IBM Tealeaf CX Configuration Manual*.

Allow Report Template Import
When enabled, the user is permitted to import report templates. When set to `Primary Group Default`, the user inherits the default setting for the primary group.

- All users can export report templates regardless of this setting.

Search Box for Online Help
Enable or disable the Online Help search box available at the top of each Portal page.

Open BBR In A New Window
When enabled, sessions loaded in Browser Based Replay are displayed in a new browser window.

Primary User Group
The primary group to use for group level settings.

Member Of
Lists all groups of which the user is currently a member.

Password

Select a user and click **Password**. Enter the new user's password twice in the **New Password** and **Confirm New Password** fields. Click **Save** to apply the change.

Force Password Change

To force the user to change his or her password, select a user and click **Force Password**. The next time the user logs in, the user is required to change the account password before using the Portal.

Clear Password History

Reset the user's password history, allowing the user to re-use passwords that might have conflicted with the Portal's password usage policy.

Assign Groups

Select a user and click **Assign Groups** to assign the user to one or more groups. To switch a user's role, remove the user from all groups, and then add a group for the user's new role.

Manage NT Users

When NT authentication is enabled in the Portal, Tealeaf administrators can use this feature to create and disable Portal accounts based on Windows domain logins detected by the Portal.

- For more information on authentication, see “Authentication” on page 329.
- NT authentication is enabled through Search Server configuration. See "Configuring the Search Server" in the *IBM Tealeaf CX Configuration Manual*.
- If desired, you can configure the Portal to automatically login and create accounts for any NT user who visits the Tealeaf Portal. See “CX Settings” on page 192.

Left pane: NT users who do not have Portal accounts **Right pane:** Current Portal accounts

- To mark an NT user for Portal account creation, select the user, and click the >> arrow.
- To disable a Portal account, select the user, and click the << arrow.

- To save these changes, click **Save**. The marked accounts are created or disabled.

View Activity

This button links to the Audit page, which lists Portal activity for the selected user. For more information see the "Audit page" section.

Delete

Deletes the selected user's profile completely, as well as any dashboards and reports created by that user.

Note:

1. When you delete a user account, any scorecard schedules or dashboards created by the user are removed. To avoid this additional removal, you may change the password on the user account instead.
2. The Admin user cannot be deleted.

Groups

Under IBM Tealeaf cxView User Administration, click **Groups** to display a list of IBM Tealeaf cxView groups.

New / Edit

Click **New** to create a new group, or select a group from the list and click **Edit** to open the group properties dialog.

User groups have the following properties. All settings except for the Group Name setting can be inherited by users who have the selected group set as their Primary User Group.

Setting Description

ID This system-generated number is used internally to identify the group.

Group Name
The name of the group.

Note: Tealeaf recommends using only alphanumeric characters when specifying names used by the Portal application. Some non-alphanumeric characters may be blocked or may cause issues for some Tealeaf components.

Default Time Zone (used in Search)
Hourly report values are displayed in the user's timezone, but the daily rollup of reporting data is controlled by the system timezone. It is best to set this value to the same as the system timezone to avoid confusion.

- This time zone setting applies primarily to the affected users' search activities. Other Portal activities utilize the Tealeaf system time zone. See "Configuring the System Timezone" in the *IBM Tealeaf CX Configuration Manual*.

Default Date Format
Select the format, such as `mm/dd/yyyy`, from the drop-down list.

Default Language
The language the Portal should display for the user group.

Default Start Page
The list of available start pages depends on the user's group and menu permissions.

Default Password Expiration
This option requires all users to change their passwords every N days. After the password expires, the portal does not let the user continue after login until a new password has been set.

Password Expiration Warning Period
The number of days before password expiration for which the system warns users that their passwords are about to expire.

Default Portal Navigation Menu
Enable or disable the group's ability to see the top navigation menu in the Portal.

Default Log User Out If Idle
When enabled, members of the group are automatically logged out if their sessions are idle for a period longer than the session timeout length.

- You can change the timeout length through Report configuration. See "Configuring the Report Server" in the *IBM Tealeaf CX Configuration Manual*.

Default Allow Report Template Import
When enabled, this setting allows all users of the group to import report templates by default. This setting can be overridden by settings to individual user accounts.

- By default, admin accounts are configured to allow report import.
- All users can export report templates regardless of this setting.

Search Box for Online Help
Enable or disable the Online Help search box available at the top of each Portal page for the group.

Open BBR In New Window
When enabled, sessions loaded in Browser Based Replay are displayed in a new browser window for the group.

Assign Users

Click **Assign Users** to add users to the selected group. The list is populated with users who are of the correct role and users who are members of absolutely no other groups.

Including users without groups allows you to transition users between roles. To assign a new role, remove all group assignments from the user, and then assign the user to another group that matches the new role.

Menu Profile

The Menu Profile for a user group determines the menu options that a group's members can reach. It also controls the reports that can be viewed on the dashboards of a group's members. For example, if a group does not have permission in the menu profile for Configure > Scorecards, its members cannot view the Scorecard Configuration page.

Note: IBM Tealeaf cxView provides a set of menu options that has been reduced from the menu options available to IBM Tealeaf cxImpact users. For more

information on available menu options in IBM Tealeaf cxImpact, see "Portal Navigation" in the *IBM Tealeaf cxImpact User Manual*.

Note: After you modify a user group's menu profile, you should verify that the users in the group can still access their default page. This page is defined at the group level in the `Default Start Page` setting and is also available for individual user configuration. Users who cannot access their default page are forced to review their settings. See "My Settings" in the *IBM Tealeaf cxImpact User Manual*.

Administration Profile

For groups granted access to the Portal Management page, the Administration Profile controls the portions of the page that they can see. The items in the list correspond to the sections on the Administration page sidebar.

Delete

Deletes the selected group. The group's members are not deleted.

Current Users

Displays a list of the currently active users of the specified role. The user's name is a link to the Audit page, which displays that user's activity.

Configuring cxView-only Users

If you have purchased licenses for IBM Tealeaf cxView-only usage, which covers scorecards and dashboards, this note describes how to configure the Tealeaf Portal for these users.

Note: To utilize IBM Tealeaf cxView-only licenses, you must be running Tealeaf version 6.1.6125 or later.

1. Login to the Portal as an administrator, with access to user and group configuration privileges.
2. In the left pane, click **IBM Tealeaf cxView User Administration**, which reveals the following three sub-menus:
   a. Users
   b. Groups
   c. Current Users
3. Click **Users**.
   - This screen allows you to configure or modify settings for individual users. Users may be incorporated directly from your Active Directory.
4. To define the settings for groups, continue to the next step. Click **Groups**.
   a. There are two default IBM Tealeaf cxView groups: `cxView User` and `cxView Admin`.
      - IBM Tealeaf cxView Admins can set dashboard options and configure scorecards.
      - IBM Tealeaf cxView Users have access to dashboards and scorecards on the Portal.
   b. If you want to have specified administrators for IBM Tealeaf cxView only, click **IBM Tealeaf cxView Admin** and select them. Otherwise, continue to configure IBM Tealeaf cxView users.
5. From the list, click **IBM Tealeaf cxView User**.

- At the bottom of this screen, you can configure basic default settings, such as date and time format, etc.

6. At the top of the screen, click **Assign Users**.
   - This screen allows you to select users from your list of accounts or from your Active Directory. These selected accounts are treated as IBM Tealeaf cxView users.
7. At the top of the screen, click **Menu Profile**.
   - Through this screen, you determine which menu options appear when users belonging to this group log onto the Tealeaf Portal. By default, IBM Tealeaf cxView users may only view reports and set basic options. You can choose to modify their permissions. For example, you can allow them to configure their own dashboards.

Expanding cxView User Permissions

IBM Tealeaf cxView user accounts can be specified to use or configure scorecards or dashboards. For example, members of the `cxView User` group of IBM Tealeaf cxView users may need to be able to use scorecards and dashboards, while members of `cxView Admin` may need to be able to use and configure both scorecards and dashboards.

IBM Tealeaf cxImpact user groups can access all of the above functionality, as well as the other features of the base IBM Tealeaf cxImpact product. For your user groups, you may wish to provide additional functionality to the IBM Tealeaf cxView users, which is available through the IBM Tealeaf cxImpact product. The solution is to turn these user accounts into IBM Tealeaf cxImpact users with additional functionality.

1. Create a IBM Tealeaf cxImpact user group with similar permissions, especially the menu profile.
   - See "CX User Administration" on page 165.
2. Add the additional features to the IBM Tealeaf cxImpact user group that you wish to enable for your IBM Tealeaf cxView users.
   - For more information on specifying the menu profile, which enables Portal features, see "Menu Profile" on page 188.
3. For each IBM Tealeaf cxView user:
   a. Remove the user from each IBM Tealeaf cxView group. When finished, the user should be a member of no groups and is not assigned to a Tealeaf product.
   b. Assign the user to the IBM Tealeaf cxImpact user group. When the assignment is complete, the user is now a IBM Tealeaf cxImpact user.
4. Repeat the above steps for each IBM Tealeaf cxView user.

Example - Enabling Drill-Down Search for cxView Users

For security reasons, IBM Tealeaf cxView users are not permitted to drill down onto scorecards and dashboards to search for the underlying sessions that contributed to the report.

However, you may decide that select IBM Tealeaf cxView users must have permission to perform drill-down searches. To enable this capability, you can create IBM Tealeaf cxImpact user accounts with the same permissions as IBM Tealeaf cxView users and the additional permissions for drill-down searches.

Note: These users are IBM Tealeaf cxImpact users at a technical level. Some of the user and group settings for IBM Tealeaf cxView do not apply to these accounts.

Note: Drill-down searches must be enabled in the individual dashboard and scorecard components to which these users have access.

- See "Configuring Scorecards" in the *IBM Tealeaf cxView User Manual*.
- See "Configuring Dashboards" in the *IBM Tealeaf cxView User Manual*.

1. Login to the Portal as an administrator.
2. From the Portal menu, select **Tealeaf** > **Portal Management**.
3. From the left navigation panel, select **IBM Tealeaf CX User Administration**.
4. Click the Groups link. See "Groups" on page 187.
5. To create a new user group, click **New**. Call it something like, `cxView - Enhanced`.
   a. Configure the following permissions:

      Setting Value

      `Default Portal Navigation Menu`
      `Disable`

   b. You may want to configure the other settings to match the group settings for the `cxView User` group. See cxView User Administration.
   c. When the new IBM Tealeaf cxImpact user group is selected, click **Menu Profile**. Assign the group the following menu permissions:

      Menu Selections

      `Dashboards`
      all

      `Analyze`
      Select any Tealeaf reports that you wish to enable drill-down for the users in the group. For IBM Tealeaf cxView, you should enable access to the `Scorecards` option at least.

      `Configure`
      `My Settings` only

      `Help` all

      - See "Menu Profile" on page 188.

   d. To save the new user, click **Save**.
6. Click the Users link.
7. To create a new user, click **New**. See "Users" on page 184.
   a. Assign the user to the `cxView - Enhanced` group.
   b. You may want to configure the other user settings to match the settings for the IBM Tealeaf cxView users. See cxView User Administration.
   c. To save the new user, click **Save**.
   d. Test the user account by logging into the Portal. Review a scorecard or dashboard where drill-down is enabled.
   e. Perform a drill-down search to see if a positive number of sessions is returned.
8. When the tests have been validated, continue creating user accounts, assigning them to the `cxView - Enhanced` group.

CX Settings

The IBM Tealeaf CX settings govern alert generation, data collection and retention, email delivery, scorecard management, search configuration, and user account management.

To access IBM Tealeaf CX settings:

1. Login to the Portal as an administrator.
2. From the Portal menu, select **Tealeaf** > **Portal Management**.
3. The Portal Management page is displayed. In the left panel, click **IBM Tealeaf CX Settings**.
4. Select the category that applies.

Note: Changes to settings saved in the Portal Management page are not applied to the Portal application until the page is refreshed through your browser.

To see information about CX configuration settings, see *IBM Tealeaf CX Configuration Manual*.

To see information resolving any issues with CX settings, see *IBM Tealeaf CX Troubleshooting Guide*.

Note: If the number of days to retain hourly event reporting data is changed, the reports might not display current data until the system normalizes.

Do not change these settings

Note: The following settings should not be modified unless you are directed to do so by your IBM Tealeaf representative.

Data collector

- `Data Aggregation`
- `Data Collection`
- `Data Collection - Batch Size`
- `Data Collector Logging Level`
- `Max Extraction table queue size`
- `Table Partitioning`
- `TDEL Buffer (minutes)` (See “Data Collector” on page 193.)

Data retention

- `Collector Data - Buffer Hours`
- `Top Mover Data (Daily)-Days Retain`
- `Top Mover Data (Hourly)-Days Retain`
- `Reporting Data (Daily)-Days Retain`
- `Reporting Data (Hourly)-Days Retain` (See “Data Retention” on page 201.)

Alerts/Thresholds

These settings control the alert generation process and the thresholds that trigger alert generation.

Table 11. Alerts/Thresholds

Setting	Description	Default
Alert Gallery Events	Specifies the list and order of alerts to display in the Alert Gallery.	N/A

Data Collector

These settings define the characteristics for the Tealeaf Data Collector, including buffer and batch sizes, intervals and time-outs, and the gathering of statistics on Tealeaf events.

- The Data Collector scans each active Canister every 5 minutes for updated information. It is a Windows service on the Report Server.

Table 12. Data Collector

Setting	Description	Default
Canister Connection - Max Concurrent Connections	The number of connections the Data Service will open to a single canister.	2
Canister Connection - Max Sequential Failures Before Logging to Event Viewer	If a data collector run fails to connect to an active canister, the event is logged in the data collector log. If the failure to connect continues to occur the number of times defined in this setting, the event is also logged in the Windows Event Viewer.	5
Canister Connection - Timeout (seconds)	The timeout for the data collector when pulling data from the Canister.	300
Data Aggregation	Set this value to Enabled to enable data aggregation of collected data for reporting purposes. **Note:** Don't change this setting unless directed by Tealeaf.	Enabled
Data Aggregation - Daily Data Processing	Describes the available settings for the frequency of daily aggregations and the date range over which they occur. **Note:** Configuring this setting to a value other than Daily through Start of Hourly Retention Period retains overlapping aggregated data in the database and may impact system performance during data aggregation. See "Data Aggregation and Retention" on page 241.	Daily through Start of Hourly Retention Period

Table 12. Data Collector (continued)

Setting	Description	Default
`Data Aggregation - Daily Data Time of Day`	This setting defines the hour of the day when the daily data aggregation run is performed, if the daily data aggregation is set to occur on a daily basis. **Note:** Tealeaf recommends configuring the daily data aggregation run to be performed during an off-peak hour.	`02:00`
`Data Aggregation - Max Concurrent Performance Threads`	The maximum number of concurrent threads that can be used for aggregating performance data.	`3`
`Data Aggregation - Max Concurrent Threads`	The maximum number of concurrent threads that can be used for aggregating base reporting data.	`4`
`Data Aggregation - Performance Daily Data Processing`	When data aggregation is performed on hourly performance data, this parameter defines the scope of the data that is aggregated at the daily level. Available options: **Note:** Configuring this setting to a value other than `Daily through Start of Hourly Retention Period` retains overlapping aggregated data in the database and may impact system performance during data aggregation. • `Hourly through Current Hour` - Performance data is aggregated to the daily level through the current hour. • `Daily through Previous Day` - Performance data is aggregated at the daily level through the previous day. • `Daily through Start of Hourly Retention Period` - Performance data is aggregated at the daily level for dates before the start of the hourly data retention period, after which the hourly data is applicable.	`Daily through Start of Hourly Retention Period`
`Data Aggregation - Performance Daily Data Time of Day`	The time of day when the data aggregation run is performed on performance data to aggregate hourly data to daily data	`4:00`
`Data Aggregation - Performance Data`	When `Enabled`, data on client performance, response times, and connection times is aggregated for reporting purposes. See "Analyzing Performance" in the *IBM Tealeaf Reporting Guide*.	`Enabled`
`Data Collection`	Set this value to `Enabled` to enable the Data Collector service to collect data from the Long Term Canister for insertion into the Tealeaf database. **Note:** Don't change this setting unless directed by Tealeaf.	`Enabled`

Table 12. Data Collector (continued)

Setting	Description	Default
Data Collection - Batch Size	How many records to extract or load at once. • Don't change this setting unless directed by Tealeaf.	2000
Data Collection - Max Concurrent	The number of canisters from which the Data Service will collect in parallel.	2
Data Collection - Maximum rows collected per canister table	If Data Collection - Limit run time is enabled, this parameter defines the maximum number of rows that the Data Collector is permitted to collect from a single Canister table on each Data Collector run. **Note:** This parameter is used in conjunction with Data Collection - Limit run time to force more frequent updates of report data, which assists the data collection process when it is far behind or struggling with a sudden spike in data volume.	2000000
Data Collection - Limit run time	When enabled, the Data Collector gathers only outstanding data with a timestamp that occurs before the start of the data collection run, instead of collecting all outstanding data. When enabled, the Data Collection - Maximum rows collected per canister table function is also enabled, and its limit is applied. **Note:** This parameter is used in conjunction with Data Collection - Maximum rows collected per canister table to force more frequent updates of report data, which assists the data collection process when it is far behind or struggling with a sudden spike in data volume.	Disabled
Data Collection Processes - Max Tries Per Staging	When an individual Data Collection run times out, the timeout setting is doubled. This process can be repeated up to the number of times defined in this setting. • The timeout is applied across all processing runs. • If the timeout is reached, the new doubled timeout setting applies until the service is restarted, at which point the timeout setting is reset to the value defined in the Database Connection - Timeout (seconds) setting.	5

Table 12. Data Collector (continued)

Setting	Description	Default
Data Collector Logging Level	Specify the logging level for the Data Collector only. **Note:** Do not change this setting unless directed by Tealeaf. Levels: • 0 - none • 1 - Error (default value) • 2 - Warning • 3 - Info • 4 - Detail • 5 - Status • 6 - Trace • 7 - All **Note:** Status level messages always appear in the log for any non-zero logging level. This value overrides the system logging level, which can be configured through TMS. See "TMS WorldView Tab" on page 262.	Error
Data Extraction - Max Table Queue Size	The maximum number of in-memory tables maintained by the Data Service. This setting is used to limit the size of the Data Service memory footprint if the speed of reading from the canister is significantly faster than writing to SQL Server. • Don't change this setting unless directed by Tealeaf.	100
Data Trimming - Canister Data	When Enabled, the un-aggregated event data is trimmed based on the configured Data Collector settings. **Note:** When this value is set to Disabled, no trimming occurs at all. The SQL Server database can grow without limit.	Enabled
Data Trimming - Canister Data Immediate Trim	When Data Trimming - Canister Data is Enabled, enabling this parameter forces the canister data tables to be immediately trimmed after the Data Collector has aggregated the data in them. **Note:** If errors are encountered during an aggregation operation, the data in with which the errors are retained. Re-aggregation and trimming is attempted during normal canister trim operations. • If Data Trimming - Canister Data is disabled, this parameter is ignored. • Don't change this setting unless directed by Tealeaf.	Enabled
Data Trimming - Interval	Determines the interval at which the Data Service trims data from the database 0 - None, 1 - Hourly, 2 - Daily, 3 - Weekly.	Hourly

Table 12. Data Collector (continued)

Setting	Description	Default
Data Trimming - Max Batch Size	The maximum number of records trimmed from the reporting or canister data tables in any single delete statement executed as part of a trimming operation.	100000
Data Trimming - Reporting Data	When Enabled, the reporting data is trimmed based on the configured Data Collector settings. **Note:** When this value is set to Disabled, no trimming occurs at all. The Reports database can grow without limit.	Enabled
Data Trimming - Statistics	When Enabled, the Tealeaf statistics data is trimmed based on the configured Data Collector settings. **Note:** When this value is set to Disabled, no trimming occurs at all. The Statistics database can grow without limit.	Enabled
Data Trimming - System	When Enabled, the user activity logs in the TL_SYSTEM database are trimmed based on the configured Data Collector settings. **Note:** When this value is set to Disabled, no trimming occurs at all. The related System database tables can grow without limit.	Enabled
Data Trimming - Time for Daily Trim	If the value for Data Trimming - Interval is Daily, then this value defines the time of day at which the reporting data is trimmed. Time is based on the Tealeaf system time zone. It should be configured for an off-peak hour.	3:00
Data Trimming - Top Movers	When Enabled, the Top Mover data is trimmed based on the configured Data Collector settings. **Note:** When this value is set to Disabled, no trimming occurs at all. Top Movers data in the database can grow without limit.	Enabled
Database Connection - Timeout (seconds)	The timeout in seconds when connecting to the database. • If the Data Collector aggregation operation times out (exceeds this setting), the setting is doubled in the next run. If it times out again, this timeout continues to be doubled until the number of times defined in the Data Collection Processes - Max Tries per Staging Table setting. The temporary extended connection timeout setting is maintained until the service is restarted, after which it reverts to the original timeout value defined for this setting.	60
Database Growth Calculation Time of Day	The time of day when the database growth report is populated with current size information • This report is available through the Portal. See "System Status" on page 8.	
5:00		
Dimension Log Aggregation	When Enabled, Tealeaf dimension values are aggregated from log entries at predefined intervals.	Enabled

Table 12. Data Collector (continued)

Setting	Description	Default
`Dimension Log Aggregation Interval`	When `Dimension Log Aggregation` is enabled, this setting defines the time interval between checks of the logs for reference values.	`Hourly`
`Dimension Log Aggregation Time`	When `Dimension Log Aggregation` is enabled and `Dimension Log Aggregation Interval` is set to `Daily`, this setting defines the 24-hour time when the review of the logs is executed.	`3:00`
`Dimension Trimming - Day of Week`	If `Dimension Trimming - Frequency` is set to `Weekly`, then this setting defines the day of the week when the trim operation is executed. • If `Dimension Trimming - Frequency` is set to `Monthly`, then this setting defines the first occurrence of the day in the month when the trim operation is executed. • See "Data Management for Dimensions" in the *IBM Tealeaf Event Manager Manual*.	`Sunday`
`Dimension Trimming - Frequency`	Set this value to how frequently the dimension trimming operation is executed: `Daily`, `Weekly`, or `Monthly`. • See "Data Management for Dimensions" in the *IBM Tealeaf Event Manager Manual*.	`Daily`
`Dimension Trimming - Time of Day`	The time of day when the dimension trimming operation occurs **Note:** Tealeaf recommends setting this value to occur during an off-peak hour, as early as possible after the end of peak usage and **after** the Scheduling Service has cycled services. • If services are cycled during a dimension trim operation, the Data Collector is forced to restart. See "Configuring the Scheduling Service" in the *IBM Tealeaf CX Configuration Manual*. • Since the Database Table Size report is updated at 2AM, changes to the table sizes are not reflected in the report until the following evening under the default setting. • See "Data Management for Dimensions" in the *IBM Tealeaf Event Manager Manual*.	`3:00`

Table 12. Data Collector (continued)

Setting	Description	Default
`Dimension Trimming - Update Fact Counts`	In addition to trimming dimension values from the dimension data, this setting enables the updating of counts of trimmed dimension values to `[others]` in all reporting data, when `Enabled`. • When `Enabled`, updating of the counts for trimmed dimension values requires making the change to every instance of the dimension value in all reporting data. Depending on the number of instances and the number of trimmed values, this process can take a few minutes to multiple hours to complete. • When `Disabled`, reporting data is not updated during dimension trimming operations. As a result, discrepancies can be introduced between the sum of event counts not filtered by the trimmed dimension and the sum of event counts filtered by the trimmed dimension. **Note:** If you do not update counts in the reporting data as part of your dimension trimming, updates for previously reviewed dimension values are not subsequently applied to the reporting data if the option is enabled at a later time. • See "Data Management for Dimensions" in the *IBM Tealeaf Event Manager Manual*.	`Enabled`
`Dimension Value Tracking - Max Concurrent Threads`	The maximum number of concurrent threads that can be spawned when timestamps for individual dimension values are being updated. • The Data Collector updates the timestamps for each value in each dimension when they are detected. During a dimension trimming operation, the Data Collector reads the timestamps associated with each value to determine the most recently occurring ones. • This setting defines the maximum number of threads that can be spawned during the timestamp updating process, which runs independently of the dimension trimming process. **Note:** Tealeaf recommends leaving this value unchanged from the default setting.	`4`
`Fact Limits - Check Interval (minutes)`	Defines the interval when the number of facts written for each event within the past hour is compared to the permitted maximum. Accepted values are `15`, `30` or `60` minutes.	`60`
`Fallback to RowByRow insert on BulkInsert Error`	When enabled and an error occurs during bulk insertion of data, the Data Collector reverts to inserting data row by row.	`Disabled`
`Send Report Schedules`	When `Enabled`, scheduled reports are delivered according to their configured settings.	`Enabled`

Table 12. Data Collector (continued)

Setting	Description	Default
`Table Partitioning`	If `Enabled`, this setting tells the Data Service that tables in the reporting database have been partitioned and need to be maintained properly. • Don't change this setting unless directed by Tealeaf.	`Disabled`
`TDEL Buffer (minutes)`	The number of minutes of already collected data to leave in the canister for recovery purposes if the system is running slowly. • Don't change this setting unless directed by Tealeaf.	`30`
`Top Movers`	When `Enabled`, Top Movers that have been configured and enabled are calculated. Hourly Top Movers are calculated once per hour, and daily Top Movers are calculated once per day.	`Enabled`
`Top Movers - Auto-calculate Daily Top Movers`	When `Enabled`, this setting forces the creation and calculation of daily Top Movers for all current and newly created events and event + dimension combinations. Switching this setting changes the status of all Top Movers in the system. You may still manually enable or disable individual Top Movers. **Note:** Enabling this setting can have significant impacts on data storage and performance. See "Data Management for Top Movers" in the *IBM Tealeaf Event Manager Manual*.	`Disabled`
`Top Movers - Auto-calculate Hourly Top Movers`	When `Enabled`, this setting forces the creation and calculation of hourly Top Movers for all current and newly created events and event + dimension combinations. Switching this setting changes the status of all Top Movers in the system. You may still manually enable or disable individual Top Movers. **Note:** Enabling this setting can have significant impacts on data storage and performance. See "Data Management for Top Movers" in the *IBM Tealeaf Event Manager Manual*.	`Disabled`
`Top Movers - Maximum data points for calculations`	The maximum number of data points required to calculate average and deviation values for event and dimension Top Movers. No more than the maximum number of data points are used in the calculation of averages and standard deviations for event and dimension Top Movers.	`16`
`Top Movers - Minimum data points for calculations`	The minimum number of data points required to calculate average and deviation values for event and dimension Top Movers. If there are too few data points, the deviation is not calculated for the period. Tealeaf recommends that you do not set this value below the default value (`4`). The minimum accepted value is `2`.	`4`

Table 12. Data Collector (continued)

Setting	Description	Default
`Top Movers - Number of Threads used for calculations`	The number of threads that the Data Collector uses when performing Top Mover hourly and daily calculations. **Note:** You can raise this setting to attempt to improve performance of Top Mover calculations. However, depending on the system load at the time of calculation, raising this setting can negatively impact system performance. See "Data Management for Top Movers" in the *IBM Tealeaf Event Manager Manual*.	`4`
`Top Movers - Time for Daily Calculation`	When `Top Movers` is enabled, this value specifies the time when deviation calculations are performed for daily Top Movers. Time is based on the Tealeaf system time zone. It should be configured for an off-peak hour.	`4:30`
`Top Movers - Calculation Mode`	This setting configures how Top Movers are computed. • `Consecutive Days` - Top Movers are calculated for consecutive days. For example, the data reported for a single Top Mover might contain entries for each day of last week and this week. • `Same Days` - Top Mover deviations are calculated from the same day from the previous weeks. For example, deviation values for Wednesday are computed using the preceding Wednesdays. • See "TEM Top Movers Tab" in the *IBM Tealeaf Event Manager Manual*.	`Consecutive Days`

Data Retention

These settings define the number of days that data is retained in the Tealeaf databases.

Note: The Data Retention settings apply to a deprecated method of generating Tealeaf alerts, which has been superseded by the Tealeaf Alert Service. These settings should not be changed. For more information on the current method, see "Configuring the Alert Service" in the *IBM Tealeaf CX Configuration Manual*.

Table 13. Data Retention

Setting	Description	Default
`Collector Data - Buffer Hours`	The number of hours of canister data to keep in the reporting database for recovery, if the system falls behind. If `Trim Collected Data Immediately` is enabled, then this setting has no effect. **Note:** Do not change this setting unless directed by Tealeaf.	`2`

Table 13. Data Retention (continued)

Setting	Description	Default
`Dimension Log Data-Days Retain`	Number of days to retain log data for dimensions. The number of days to retain data is relative to the last instance of a reference value being detected in the traffic. When a new instance of the value has been detected, the number of days for that dimensional value is reset to zero.	`30`
`Dimension Trimming - # of Most Recent Values to Keep`	The maximum number of values to retain for any dimension at any time. If the number of retained values is above this limit when the Data Collector executes, the most recent values as determined by this setting are retained, and all older values are trimmed. See "Data Management for Dimensions" in the *IBM Tealeaf Event Manager Manual*.	`750,000`
`Dimension Log Capture Days`	Defines the number of days to log a dimension. Logging of the dimension is disabled after the defined number of days have expired. Dimension logging may be enabled and disabled per dimension in the dimension's definition. See "TEM Dimensions Tab" in the *IBM Tealeaf Event Manager Manual*. Dimension logging may be enabled and disabled per dimension in the dimension's definition. See "TEM Dimensions Tab" in the *IBM Tealeaf Event Manager Manual*.	`14`
`Reporting Data (daily)-Days Retain`	Controls the retention of daily-level reporting data. **Note:** Do not change this setting unless directed by Tealeaf.	`365`
`Reporting Data (hourly)-Days Retain`	Controls the retention of hourly-level reporting data. **Note:** Do not change this setting unless directed by Tealeaf.	`30`
`System Processing Stats-Days Retain`	Number of days to keep the hourly statistics data.	`30`

Table 13. Data Retention (continued)

Setting	Description	Default
`Top Mover Data (daily)-Days Retain`	Controls the retention of daily-level Top Mover data. **Note:** 1. Do not change this setting unless directed by Tealeaf. 2. This setting must be configured in conjunction with the `Top Movers - Minimum data points for calculations` and `Top Movers - Maximum data points for calculations` settings, which define the minimum and maximum number of data points to retain. Since each retained data point corresponds to a week of data, your retention period settings must be configured to be greater than the minimum and less than or equal to the maximum data points. Multiply the data point settings by 7 to make the comparison to this setting.	`365`
`Top Mover Data (hourly)-Days Retain`	Controls the retention of hourly-level Top Mover data. **Note:** 1. Do not change this setting unless directed by Tealeaf. 2. This setting must be configured in conjunction with the `Top Movers - Minimum data points for calculations` and `Top Movers - Maximum data points for calculations` settings, which define the minimum and maximum number of data points to retain. Since each retained data point corresponds to a week of data, your retention period settings must be configured to be greater than the minimum and less than or equal to the maximum data points. Multiply the data point settings by 7 to make the comparison to this setting.	`120`
`User Activity Log Data-Days Retain`	Number of days to keep the Portal user activity history data.	`90`

Search

IBM Tealeaf CX search settings control how session searches are configured and retrieved. Searches may return a high number of sessions, so some settings can be used to limit the volume of retrieved data.

Table 14. Search

Setting	Description	Default
BBR Annotation Options	Administrators can configure the list of pre-selected annotation options available in Browser-Based Replay. Select this option and scroll down to review and update the list. Users can also create their own annotations through BBR. For more information, see "Tracking Interactions through BBR" in the *IBM Tealeaf cxImpact User Manual*.	N/A
Completed Max Results Limit	Max number of sessions to return for a completed session search. See "Evaluating Search Results Limits" on page 206.	100
Create Session Segment for Completed Searches	Global setting to enable session segments for searches for completed sessions. This can be disabled at a user level by the user setting.	Enabled
Default Initial Session Segment Days To Live	The default lifespan for session segments. Individual session segments can have expiration postponed from the Session Segments page.	1
Default Session List Template	The session list template to use for drilldown searches for IBM Tealeaf CX users. Searches executed from the Portal search page utilize the assigned or selected session list template. See "Searching Session Data" in the *IBM Tealeaf cxImpact User Manual*.	<Default>
Display Event Info Tooltips	Toggle display of tooltips for events.	Enabled
Display Passive Capture Info	Show/hide PCA info when drilling down to a specific page of a session.	Enabled
Event Image Display Max	The maximum number of events to display in the Event List column on the Session List.	20
Event Selector Dimensions	The initial size for the Event Selector.	2x1

Table 14. Search (continued)

Setting	Description	Default
Hit Detail Enabled	Enables the Page Detail view in search results (**Session List** > **Session Detail** > **Page Detail**).	Enabled
Live Max Results Limit	The maximum number of active sessions to return from an active search. See "Evaluating Search Results Limits" on page 206.	100
Live Search Time-Out (seconds)	The maximum time to allow for an active search to run.	600
Open Browser Replay in Featureless Window	Enables the removal of toolbars and buttons from the Browser Based Replay window, which will force a new window and not a new tab in the browser.	Enabled
Require At Least One Populated Search Field	Require at least one field among the search options to be populated before a search can be executed	Disabled
Results Per Page	The page size of the session list grid in search results.	20
Session Detail Path Variable	The field to retrieve the URL for each page in session detail.	env/URL
Session Segment Max Sessions Limit	The maximum number of sessions to use in a session segment. **Note:** Increasing this setting may impact system performance. Some Tealeaf components remove session information when segments are created containing more than 16,384 sessions. Tealeaf recommends retaining the default setting and performing narrower searches. See "Evaluating Search Results Limits" on page 206. **Note:** If you have upgraded from Release 7.x or earlier and are experiencing issues with drill-down searches not working through the Report Builder, you may need to adjust this setting. See "Troubleshooting - Reporting" in the *IBM Tealeaf Troubleshooting Guide*.	1000
Session Segment Sharing	Enables session segments to be shared between users.	Enabled

Table 14. Search (continued)

Setting	Description	Default
`Visitor Search Default Top N Report Data Rows Returned`	Maximum number of rows of sessions that can be returned in a session list through Visitor search. This setting is used to limit unnecessarily large data transfers in Visitor search results. See "Evaluating Search Results Limits." **Note:** This setting applies only if IBM Tealeaf cxResults has been licensed and installed. IBM Tealeaf cxResults is no longer available as a newly licensed product as of Release 8.7. Customers that licensed IBM Tealeaf cxResults in Release 8.6 and earlier may continue to use and receive support for the product in Release 8.7 and later. For more information, please contact Tealeaf Customer Support . • See "Searching for Visitors" in the *IBM Tealeaf cxResults User Manual*.	`1000`
`Visitor Search Template`	The Visitor search template to use for visitor searches. **Note:** This setting applies only if IBM Tealeaf cxResults has been licensed and installed. IBM Tealeaf cxResults is no longer available as a newly licensed product as of Release 8.7. Customers that licensed IBM Tealeaf cxResults in Release 8.6 and earlier may continue to use and receive support for the product in Release 8.7 and later. For more information, please contact Tealeaf Customer Support .	`<Default Completed>`

Evaluating Search Results Limits

In the Search settings, you can define the maximum number of results returned when searching for active or completed sessions.

- If needed, you can define the properties to have different values. Since active sessions are not indexed, you may want to set that property to a lower value than the value for completed sessions.

Additionally, if you are using IBM Tealeaf cxResults, an additional limit is imposed on the number of results returned in a Visitor search: `Visitor Search Default Top N Report Data Rows Returned`.

These limit settings, however, can be overridden by values defined in individual search templates.

Note: When evaluating the maximum number of search results to return, the lower value of the limits defined in the above search settings or in the search template is applied.

- For more information on the search template limits, see "Configuring Search Templates" in the *IBM Tealeaf cxImpact Administration Manual*.

Users

These IBM Tealeaf CX settings control creation of user accounts, password policies, and default user settings.

Table 15. Users

Setting	Description	Default
Automatic NT User Login	If NT Auth is enabled, this setting allows visitors to the Tealeaf Portal to be automatically logged in using their Windows NT domain credentials, by passing the Portal login screen.	Enabled
Automatic User Creation	If NT Auth is enabled, this setting allows for new users to have Portal accounts automatically created on first login.	Enabled
Automatic User Creation Settings Prompt	Enables first-time users to specify their user settings immediately upon login.	Disabled
Automatic User Creation Type	The type of user account to create when a user is auto-created.	Portal
Display Splash Screen on User Login	When enabled, users see the Portal Announcements window when they login.	Enabled
Login Page Language (Global)	The language used on the global login page. After the user logs in, the user's preferred language is known.	English (US)
Minimum Password Length	Minimum password length for user passwords. • Portal passwords have a maximum length of 32 characters.	3
New User Default Group	The user group to assign to automatically created users, one for each role: IBM Tealeaf CX, IBM Tealeaf cxReveal, or IBM Tealeaf cxView.	Portal User
New User Default Page	The default start page to assign to automatically created users.	1
New User Default Time Zone	The default timezone to assign automatically created users.	System time zone
New User Group Default Language	The default language automatically assigned to newly created users.	English (US)

Table 15. Users (continued)

Setting	Description	Default
Previous Password History (Count)	Number of unique passwords a user must have before reusing a password.	Disabled
Previous Password History (Days)	Number of days that must pass before a user can reuse a password.	-1
Require Strong Passwords	Require passwords to meet the 'strong password' requirements.	Disabled

Integration

These settings are used to define the Segment Builder, which is a service that can create a session segment from an external data source.

Note: These settings apply to IBM Tealeaf cxConnect for Data Analysis, a separately licensable product for the IBM Tealeaf CX system. please contact your IBM Tealeaf representative.

Table 16. Integration

Setting	Description	Default
Apply MD5 to value	Apply MD5 hash to session identifier column.	Disabled
Column for data field	The column number for the session identifier.	1
Delimiter character for data records	The delimiter character is used to separate columns in the source data file.	,
Digital Analytics Explore - Client ID	The client ID name is returned in the exported data file that is used by IBM® Digital Analytics.	
Digital Analytics Explore - FTP Export Password	The FTP export password is used to log on to the FTP server where the Tealeaf data file is stored.	
Digital Analytics Explore - FTP Export Username	The FTP Username is used to log on to the FTP server where the Tealeaf data file is stored.	
Digital Analytics Explore - FTP Hostname	The FTP host name identifies the FTP server where the Tealeaf data file is stored.	

Table 16. Integration (continued)

Setting	Description	Default
Digital Analytics Explore - FTP Import Password	The FTP import password is used to log on to the FTP server where the IBM Digital Analytics data file is stored for use by Tealeaf.	
Digital Analytics Explore - FTP Import Username	The FTP import user name is used log on to the FTP server where the IBM Digital Analytics data file is stored for use by Tealeaf.	
FTP Port	The FTP port is the port number that is used for FTP file transfers.	`21`
Digital Analytics Explore - FTP Use SSL	Enable or disable SSL to enable or disable secure FTP file transfers using SSL.	`Disabled`
Digital Analytics Explore - Hostname	The host name of the IBM Digital Analytics server.	`explore.coremetrics.com`
Directory for processed files	The location where processed files are stored. If a drive letter is not specified, then the path is relative to the Tealeaf install directory.	`Integration\Completed`
Directory for segment files	The location where the source files are stored. If a drive letter is not specified, then the path is relative to the Tealeaf install directory.	`Integration\Segments`
Enable Segment Directory Watcher	When this option is enabled, the Segment Builder service polls the upload directory at regular intervals for files to use to build new segments. If the segment files are large (> 1 MB), this option should be disabled. The service checks for new files every 5 minutes when this option is disabled. • See "Segment Builder Service" in the *IBM Tealeaf cxConnect for Web Analytics Administration Manual*.	`Enabled`

Table 16. Integration (continued)

Setting	Description	Default
Integration ID prefix	If this option is enabled, the search value is parsed from the data file record using the prefix string defined in the `Integration ID prefix` setting as the start and through the delimiter or an end of line as the end of the string. • If this option is disabled, then the parser iterates through each line using a defined column as the integration key. • When building a segment, this value is used with the `Delimiter character for data records` setting to parse the search value from the data. This prefix is used when the data may have multiple values that are concatenated into one field in the data file. • If the search term is clearly defined in a column of the data file, this configuration option is not necessary.	`TLTIID_`
Name for Session Segment	The name for the generated session segment when viewed through the Tealeaf Portal.	`Segment Builder`
Search date span (days)	The number of preceding days of session data that are searched to build the session segment from the source files.	`3`
Search term for data field	The technical term used in the search to build the session segment. **Note:** Do not change unless you are familiar with the Tealeaf search syntax.	`tltstssesnid`
Session Segment Owner	The default owner for the session segment.	`ADMIN`
Share Segment	Set this option to `Enabled` to share segments created by the Segment Builder service with all Tealeaf users after they are created.	`Enabled`
Use prefix value to parse integration ID field	If this option is enabled, the integration parser removes from each line of data from the first instance of the prefix string defined in `Integration ID prefix` through the delimiter or an end of line. **Note:** If this option is disabled, then the parser iterates through each line using a defined column as the integration key.	`Enabled`

Miscellaneous

Table 17. Miscellaneous

Setting	Description	Default
`Data Service Dimension Value Paging Limit`	This value defines the maximum number of dimension values per page that can be retrieved through the Data Service. Do not change this setting unless you are experiencing performance problems with the Data Service. See "Configuring the Tealeaf Data Service" in the *IBM Tealeaf CX Configuration Manual*.	`10,000`
`Data Service SQL Trace`	When enabled, this setting causes the Data Service to log SQL queries from the Data Layer to disk. These queries are stored in `Logs\TLDataService-SqlTrace-YYYYMMDD`. **Note:** SQL trace statements can quickly grow the log file. Do not enable this option unless you are actively troubleshooting an issue. Disable it when the issue has been resolved. Do not change this setting unless you are experiencing performance problems with the Data Service. See "Configuring the Tealeaf Data Service" in the *IBM Tealeaf CX Configuration Manual*.	`Disabled`
`Excel Export Extension`	By default, Tealeaf exports Excel reports in an XML format that can be read by Microsoft Excel and uses the `xml` extension for the filename. However, `xml` files may be mapped to a different application. You may change this setting to export the reports with a different filename extension so that they can be automatically opened in Microsoft Excel. Typically, this value is changed to `xls`. When the file is loaded in Excel, users may receive an error message, which can be ignored. If desired, you can configure your local system to load all XML documents through the appropriate Microsoft Office application. See "Troubleshooting - Portal" in the *IBM Tealeaf Troubleshooting Guide*.	`xml`
`Extended Duplicate Canister Server Detection`	When enabled, the Portal checks to verify that any changes to a Canister Server definition in the Portal do not result in a conflict with another Canister Server's host name, IP address, or full qualified domain name. Duplicate Canister Server references can cause data extraction problems. **Note:** The above setting is enabled by default. It should be disabled only in environments where network infrastructure makes DNS queries complete very slowly. If this setting is disabled, then you should manually verify that there are no duplicate Canister Servers in the system. See “Managing Tealeaf Servers” on page 115.	`Enabled`

When properly configured, the following two settings enable Portal users to automatically login to Tealeaf Online Help where official product documentation is maintained and updated.

Table 18. Miscellaneous

Setting	Description	Default
Online Help Customer ID	The username to access Tealeaf Online Help.	N/A
Online Help Password	The password to access Tealeaf Online Help.	N/A

Note: Do not change the Customer Name and Password for Online Help without first informing http://support.tealeaf.com/.

These settings can be used to specify the individual at your enterprise to whom inquiries about the Tealeaf solution should be directed. This individual should have administrator access to the Tealeaf solution.

- The values for these settings are displayed in the About dialog available through the Help menu in the Portal.

Table 19. Miscellaneous

Setting	Description	Default
Prohibit Portal from Running Inside a Browser Frame	When Enabled, the Portal is always launched within the main browser window.	Disabled
Report Builder Database Connection Timeout (seconds)	This setting defines the number of seconds that a query initiated through the Report Builder in the Portal is allowed to run without a response, after which the query times out.	600
Show Disabled Report Groups Alert Message	When Enabled, a warning message is displayed across the top of the Portal whenever one or more report groups has been disabled due to too many facts being recorded within a single hour. For more information, see "Portal State Messages" in the *IBM Tealeaf cxImpact User Manual*.	Enabled
Show Spool Alert Message	When Enabled, a warning message is displayed across the top of the Portal whenever one or more Canisters is spooling data to disk. See "Portal State Messages" in the *IBM Tealeaf cxImpact User Manual*.	Enabled
System Authentication Verification Announcement	When Enabled, the Portal displays a Portal Announcement message that the Portal is improperly configured and includes the list of misconfigured or unresponsive servers. See “Portal Announcements” on page 226.	Enabled
Tealeaf Administrator E-mail	Email account of the Tealeaf administration contact. **Note:** This email address is used by multiple Portal features and should always be specified. If it is not specified, users may not be immediately updated of forced changes in behavior, including the automated disabling of some features.	N/A
Tealeaf Administrator Name	The user-friendly name of the Tealeaf administration contact.	N/A
Tealeaf Administrator Phone	The phone number of the Tealeaf administration contact.	N/A

cxReveal Settings

These settings govern the behavior of IBM Tealeaf cxReveal, which is used to search, analyze, and replay visitor sessions.

To accessIBM TealeafcxReveal settings:

1. Login to the Portal as an administrator.
2. From the Portal menu, select **Tealeaf** > **Portal Management**.
3. The Portal Management page is displayed. In the left panel, click **IBM Tealeaf cxReveal Settings**.
4. Select the category that applies.

Note: Changes to settings saved in the Portal Management page are not applied to the Portal application until the page is refreshed through your browser.

Search

cxReveal search settings control how session searches are configured and retrieved. Searches may return a high number of sessions, so some settings can be used to limit the volume of retrieved data.

Table 20. Search

Setting	Description	Default
BBR Annotation Options	Administrators can configure the list of pre-selected annotation options available in Browser-Based Replay. Select this option and scroll down to review and update the list. Users can also create their own annotations through BBR • See "Tracking Interactions through BBR" in the *IBM Tealeaf cxImpact User Manual*.	N/A
Completed Max Results Limit	Max number of sessions to return for a completed session search. See "CX Settings" in the *IBM Tealeaf cxImpact Administration Manual*.	100
Default Session List Template	The session list template to use for drill-down searches for IBM Tealeaf cxReveal users. • Searches executed from the Portal search page utilize the assigned or selected session list template. See "Searching Session Data" in the *IBM Tealeaf cxImpact User Manual*.	<Default>
Display Event Info Tooltips	Toggle display of tooltips for events.	Enabled
Display Passive Capture Info	Show/hide PCA information when drilling down to a specific page of a session.	Enabled
Event Image Display Max	The maximum number of events to display in the Event List column on the Session List.	20
Event Selector Dimensions	The initial size for the Event Selector.	2x1
Hit Detail Enabled	Enables the Page Detail view in search results (**Session List > Session Detail > Page Detail**).	Enabled

Table 20. Search (continued)

Setting	Description	Default
Live Max Results Limit	The maximum number of active sessions to return from an active search. See "CX Settings" in the *IBM Tealeaf cxImpact Administration Manual*.	100
Log Search Database Searches	When enabled, searches of the IBM Tealeaf cxReveal database are logged. • This setting is present only if the IBM Tealeaf cxReveal database has been installed. • The number of days that logs are retained is defined by Search Database Log - Days Retain.	Disabled
Open Browser Replay In Featureless Window	Enables the removal of toolbars and buttons from the Browser-Based Replay window, which will force a new window and not a new tab in the browser.	Enabled
Require At Least One Populated Search Field	When enabled, searches executed through IBM Tealeaf cxReveal must contain at least one populated search field to limit their scope.	Disabled
Results Per Page	The page size of the session list grid in search results.	20
Search Database Log - Days Retain	When Log Search Database Searches is enabled, this setting defines the number of days for which logs of IBM Tealeaf cxReveal database searches are retained. • This setting is present only if the IBM Tealeaf cxReveal database has been installed.	7
Search Database Max Sessions Returned	Specifies the maximum number of sessions that can be returned from a database search. • This setting defines the volume of search data returned from the IBM Tealeaf cxReveal Search database. Other settings that limit the number of displayed sessions, such as Completed Max Results Limit, may be applied afterward, further limiting the number of search results.	100
Search Database Search Timeout (seconds)	This setting defines the maximum amount of time in seconds that a search is permitted to run before it times out. • This setting applies only if you have installed the IBM Tealeaf cxReveal database. See "cxReveal Installation" in the *IBM Tealeaf cxReveal Administration Manual*.	60
Search Database Session Attribute Search Tables - Days Retain	Number of days of session attribute data to retain in the IBM Tealeaf cxReveal database tables • This setting applies only if you have installed the IBM Tealeaf cxReveal database. See "cxReveal Installation" in the *IBM Tealeaf cxReveal Administration Manual*.	30

Table 20. Search (continued)

Setting	Description	Default
Search Database Session Timeout (minutes)	For IBM Tealeaf cxReveal, this setting defines the number of minutes that are permitted to elapse without additional activity, after which the session is automatically timed out. See "Configuring Session Attribute Search" in the *IBM Tealeaf cxReveal Administration Manual*. • This setting applies only if you have installed the IBM Tealeaf cxReveal database. See "cxReveal Installation" in the *IBM Tealeaf cxReveal Administration Manual*.	60
Session Detail Path Variable	The field to retrieve the URL for each page in session detail.	env/URL
Use Canister Data for Database Search Results	When enabled, searches executed against the IBM Tealeaf cxReveal database to retrieve session attribute information also gather session information from the Canisters where the matching sessions are stored. **Note:** Tealeaf recommends keeping this setting enabled, which is the default configuration. However, if performance degrades significantly for searches against the IBM Tealeaf cxReveal database, you may try disabling this setting. See "Configuring Session Attribute Search" in the *IBM Tealeaf cxReveal Administration Manual*.	Enabled

Users

These IBM Tealeaf cxReveal settings control creation of user accounts, password policies, and default user settings.

Table 21. Users

Setting	Description	Default
Automatic User Creation	If NT Auth is enabled, this setting allows for new users to have Portal accounts automatically created on first login.	Enabled
Automatic User Creation Settings Prompt	Enables first-time users to specify their user settings immediately upon login.	Disabled
Login Page Language (Global)	The language used on the global login page. After the user logs in, the user's preferred language is known.	English (US)
Minimum Password Length	Minimum password length for user passwords. Portal passwords can have a maximum length of 32 characters.	3
New User Default Page	The default start page to assign to automatically created users.	Report Builder

Table 21. Users (continued)

Setting	Description	Default
New User Default Time Zone (used in search)	The default timezone to assign automatically created users. This setting is used for display in search. By default, this value is set to the Tealeaf system time zone. See "Configuring the System Timezone" in the *IBM Tealeaf CX Configuration Manual*.	System time zone
Previous Password History (Count)	Number of unique passwords a user must have before reusing a password.	Disabled
Previous Password History (Days)	Number of days that must pass before a user can reuse a password.	Disabled
Require Strong Passwords	Require passwords to meet the 'strong password' requirements.	Disabled
User Default Group	The user group to assign to automatically created users, one for each role: IBM Tealeaf cxImpact, IBM Tealeaf cxReveal, or IBM Tealeaf cxView.	cxReveal User

cxView Settings

These settings control how IBM Tealeaf cxView is configured.

To accessIBM TealeafcxView settings:

1. Login to the Portal as an administrator.
2. From the Portal menu, select **Tealeaf** > **Portal Management**.
3. The Portal Management page is displayed. In the left panel, click **IBM Tealeaf cxView Settings**.
4. Select the category that applies.

Note: Changes to settings saved in the Portal Management page are not applied to the Portal application until the page is refreshed through your browser.

Scorecards

Use these settings to configure schedule retention, dashboard distribution, and scorecard data retention.

Table 22. Scorecards

Setting	Description	Default
Data Trimming	When Enabled, scorecard data is trimmed according to the configured trim settings. **Note:** This setting does not affect data trimming for base IBM Tealeaf cxImpact data, which is stored and managed independently of scorecard data. For more information on those settings, see cxImpact Settings.	Enabled
Schedule History - Days Retain	The number of days to retain the history of scheduled scorecard runs.	30
Scorecard Data - Days Retain	The number of days to keep Day scorecard calculations.	365
Scorecard Data - Months Retain	The number of months to keep Month scorecard calculations.	24
Scorecard Data - Weeks Retain	The number of weeks to keep Week scorecard calculations.	104
Scorecard Email - Attach HTML	Attach the scorecard HTML as an attachment to the e-mail.	Disabled
Scorecard Update Frequency	Frequency in minutes when scorecards are updated with data from the database. Default value is 60 minutes. You may also set the frequency to 5 minutes, which is the rate at which the Data Collector gathers data from the Processing Servers and aggregates and inserts it into the reporting database. • Tealeaf recommends leaving this value at 60. Updating scorecards every five minutes can impact performance.	60
Temporary Directory	The temp directory for any scheduled reports that generate PDFs, relative to the directory in which the Tealeaf Data Service is running. • The default directory is <Tealeaf_install_directory>\Reporting. • The default value for this setting is the same directory.	.

Users

These IBM Tealeaf cxView settings control creation of user accounts, password policies, and default user settings.

Table 23. Users

Setting	Description	Default
Automatic User Creation	If NT Authentication is enabled, this option allows new users to have Portal accounts automatically created on first login.	Enabled
Automatic User Creation Settings Prompt	Enables first-time users to specify their user settings immediately upon login.	Disabled
Event Selector Dimensions	The initial size for the Event Selector.	2x1
Login Page Language (Global)	The language used on the global login page. After the user logs in, the user's preferred language is known.	English (US)
Minimum Password Length	Minimum password length for user passwords. • Portal passwords can have a maximum length of 32 characters.	3
New User Default Page	The default start page to assign to automatically created users.	Default Dashboard
New User Default Time Zone (used in Search	The default timezone to assign automatically created users. This setting is used for display in search. • By default, this value is set to the Tealeaf system time zone. See "Configuring the System Timezone" in the *IBM Tealeaf CX Configuration Manual*.	System time zone
Previous Password History (Count)	Number of unique passwords a user must have before reusing a password.	Disabled
Previous Password History (Days)	Number of days that must pass before a user can reuse a password.	Disabled
Require Strong Passwords	Require passwords to meet the 'strong password' requirements.	Disabled
User Default Group	The user group to assign to automatically created users, one for each role: IBM Tealeaf cxImpact, IBM Tealeaf cxReveal, or IBM Tealeaf cxView.	cxView User

cxResults Settings

The IBM Tealeaf cxResults settings govern user-configurable options for the IBM Tealeaf cxResults (VISITORS) databases and the server that hosts it.

Note: IBM Tealeaf cxResults is no longer available as a newly licensed product as of Release 8.7. Customers that licensed IBM Tealeaf cxResults in Release 8.6 and earlier may continue to use and receive support for the product in Release 8.7 and later. For more information, please contact Tealeaf Customer Support .

Note: Changes to settings saved in the Portal Management page are not applied to the Portal application until the page is refreshed through your browser.

Do not change these settings

Note: The following settings should not be modified unless you have been directed to do so by Tealeaf.

Visitor extract

- Event Table Commit Size
- Insert to DB
- Session Table Commit Size (See "Visitor Extract.")

Visitor Extract

The Visitor Extract settings pertain to the extraction process that pulls visitor data from the Processing server(s) and loads it into the IBM Tealeaf cxResults databases.

Table 24. Visitor Extract

Setting	Description	Default
Database Connection Timeout (seconds)	The length of time in seconds that a statement in the connection is allowed to execute.	900
Event Table Commit Size	The number of rows in the event table that are kept in memory before the data is committed to the database. • Don't change this setting unless directed by Tealeaf.	1000
Extractor queue empty wait interval (milliseconds)	The number of milliseconds for which an empty queue is waited to be filled.	200
Extractor queue max	The maximum number of sessions that can be queued for processing by any available thread.	4
Extractor threads	The number of threads the Visitor extract process uses for loading data into the staging database. • By default, this value is set to 0, in which case the number of threads to use is automatically calculated as the lesser of the number of Tealeaf Canisters and the number of available CPUs by the Visitor Database Extractor. If these values are not available, the number of threads is set to 1 and recalculated on next extract. **Note:** Tealeaf recommends leaving this value at 0. Setting a thread count above the calculated value does not improve system performance.	0

Table 24. Visitor Extract (continued)

Setting	Description	Default
Include Session Attributes in Visitor Extract	Toggle to enable loading of custom session attributes into the IBM Tealeaf cxResults database.	Enabled
Insert to DB	If enabled, data records are inserted into the database. If disabled, data records are discarded. **Note:** This setting is for testing purposes only. You can use this setting to distinguish whether issues are being caused by data retrieval or data storage. • Don't change this setting unless directed by Tealeaf.	Enabled
Maximum days to trim at once from Reporting tables	Maximum number of days of data to trim in one pass from the tables of the Visitor Reporting database.	2
Populate Event TextFound column in Visitor Extract	When Enabled, the TextFound value for each configured event is stored in the Visitor database. **Note:** This setting must be enabled in order to make event values available for Visitor Search. If event values are not stored, there can be significant savings of storage space.	Enabled
Processed sessions reporting interval	The interval at which sessions processed in the Visitor database are reported.	1000
Records per insert from Staging to Reporting tables	The maximum number of records to add per insert command executed from the staging tables.	20000
Records per trim from Reporting tables	The maximum number of records to remove per trim command executed on the reporting tables.	100000
Reporting Visitor Data-Days Retain	This value is the number of days of data retained for reporting purposes.	30
Session Attribute - LoginId	The session attribute that contains the Login ID information. This session attribute is used as a reference variable in IBM Tealeaf cxResults reporting. • In new installations of Release 8.0 and later, this setting is unlikely to need to be changed. See "Post-Installation Tasks" in the *IBM Tealeaf CX Installation Manual*. **Note:** In Release 7.2 and earlier, this data was stored in the Login ID reference variable. During upgrade, the Login ID variable is automatically assigned to Custom Variable 00. In most Tealeaf solutions, no further configuration is necessary. You may change this session attribute if you know that your solution has been configured to use a different session attribute to store the login identifier for your web application.	Login ID

Table 24. Visitor Extract (continued)

Setting	Description	Default
Session Attribute - Userdef1	The session attribute that contains the UserDef1 information. This session attribute is used as a reference variable in IBM Tealeaf cxResults reporting. • In new installations of Release 8.0 and later, this setting is unlikely to need to be changed. See "Post-Installation Tasks" in the *IBM Tealeaf CX Installation Manual*. **Note:** In Release 7.2 and earlier, this data was stored in the `UserDef 1` reference variable. During upgrade, this variable is automatically assigned to `Custom Variable 01`. In most Tealeaf solutions, no further configuration is necessary. You may change this session attribute if you know that your solution has been configured to use a different session attribute to store the login identifier for your web application.	`Custom Variable 01`
Session Attribute - Userdef2	The session attribute that contains the UserDef2 information. This session attribute is used as a reference variable in IBM Tealeaf cxResults reporting. • In new installations of Release 8.0 and later, this setting is unlikely to need to be changed. See "Post-Installation Tasks" in the *IBM Tealeaf CX Installation Manual*. **Note:** In Release 7.2 and earlier, this data was stored in the `UserDef 2` reference variable. During upgrade, this variable is automatically assigned to `Custom Variable 02`. In most Tealeaf solutions, no further configuration is necessary. You may change this session attribute if you know that your solution has been configured to use a different session attribute to store the login identifier for your web application.	`Custom Variable 02`
Session Attribute - Userdef3	The session attribute that contains the UserDef3 information. This session attribute is used as a reference variable in IBM Tealeaf cxResults reporting. • In new installations of Release 8.0 and later, this setting is unlikely to need to be changed. See "Post-Installation Tasks" in the *IBM Tealeaf CX Installation Manual*. **Note:** In Release 7.2 and earlier, this data was stored in the `UserDef 3` reference variable. During upgrade, this variable is automatically assigned to `Custom Variable 03`. In most Tealeaf solutions, no further configuration is necessary. You may change this session attribute if you know that your solution has been configured to use a different session attribute to store the login identifier for your web application.	`Custom Variable 03`

Table 24. Visitor Extract (continued)

Setting	Description	Default
Session Attribute - Userdef4	The session attribute that contains the UserDef4 information. This session attribute is used as a reference variable in IBM Tealeaf cxResults reporting. • In new installations of Release 8.0 and later, this setting is unlikely to need to be changed. See "Post-Installation Tasks" in the *IBM Tealeaf CX Installation Manual*. **Note:** In Release 7.2 and earlier, this data was stored in the `UserDef 4` reference variable. During upgrade, this variable is automatically assigned to `Custom Variable 04`. In most Tealeaf solutions, no further configuration is necessary. You may change this session attribute if you know that your solution has been configured to use a different session attribute to store the login identifier for your web application.	`Custom Variable 04`
Session Table Commit Size	The number of rows in the session table that are kept in memory before the data is committed to the database. • Don't change this setting unless directed by Tealeaf.	`100`

Visitor Server

These settings apply to the reporting database of the Visitors server.

Table 25. Visitor Server

Setting	Description	Default
Database Segmentation Update Interval	When Data Segmentation is enabled through Search Server, this setting defines how frequently IBM Tealeaf cxResults queries Search Server for updates to event filters.	`Hourly`
Database Connection Timeout(seconds)	The length of time in seconds that a statement in the connection is allowed to execute.	`600`
Enable Data Partitioning	When `Enabled`, data in the Visitor database is automatically partitioned.	`Enabled`
Maximum Parallel Segment Queries	The maximum number of queries that can be executed at the same time.	`4`
Maximum Parallel Text Searches	The maximum number of text searches that can be executed at the same time.	`4`

Table 25. Visitor Server (continued)

Setting	Description	Default
Maximum Text Search Interval (minutes)	The maximum time interval for a text search. Queries that include text terms and have a time span greater than this value will be broken up into multiple smaller text searches. For example, if this setting is set to 60, a text search in IBM Tealeaf cxResults that spans two hours and one minute is broken up into three separate queries. • The default value is 60 (one hour). **Note:** This setting defines the primary upper limit in terms of size for text-based searches in IBM Tealeaf cxResults. Depending on the hardware of your system and your site traffic volume, increasing this limit may cause issues. High-volume sites that set this value too high may generate search failures.	60
Text Search Max Result	Specifies the maximum number of results returned in a text-based search. This limit is intended to prevent runaway searches. **Note:** Increasing this value can cause the system to run out of memory. • This limit applies to all IBM Tealeaf cxResults searches that do not search for event data.	10000000

RS Extractor Settings

This section contains the settings for the RSE Extractor Service (RSE), which is used to analyze session segments.

To access RS Extractor settings:

1. Login to the Portal as an administrator.
2. From the Portal menu, select **Tealeaf** > **Portal Management**.
3. The Portal Management page is displayed. In the left panel, click **RS Extractor Settings**.
4. Select the category that applies.

Note: Changes to settings saved in the Portal Management page are not applied to the Portal application until the page is refreshed through your browser.

Don't Change These Settings

RS Extractor Settings

- **Time in seconds to verify the JobManager is running**
- **Time interval in seconds between checks of system configuration settings**
- **Activates Profiling**
- **Maximum number of permitted search server connections**
- **Maximum number of sessions queued in memory for processing by a job**
- **Number of unwritten data rows to accumulate before writing**

Service Settings

Setting Description

Time in seconds to verify the JobManager is running
How often the RSE service performs a status ping.

- Don't change this setting unless directed by Tealeaf.

Time interval in seconds between checks of system configuration settings
How often the RSE service reloads its configuration information.

- Don't change this setting unless directed by Tealeaf.

Time interval in seconds between checks for Resultsets to remove
How long the RSE service waits between checks for expired session segments.

Enable Resultset expiration warnings
If set to True, emails are sent prior to session segments expiring.

Time in seconds before sending expiration warnings
How long the RSE service waits between checks for expiration warnings.

Enable Job completion notifications
If set to True, emails are sent when session segment jobs complete.

Enable Profiling
If set to True, a detailed profile is provided for the RSE service performance.

- Don't change this setting unless directed by Tealeaf.

Email format
Set to plain text or HTML.

Enable Datastore refresh
Set to True, if the RSE service should clear its memory cache.

Logging level for job history
The level of logging the RSE service does.

Maximum number of sessions in a job
The maximum number of sessions a session segment can contain.

Maximum number of resultset sets
The maximum number of session segments that are allowed.

Number of days to keep resultsets
The number of days that session segments are retained.

Enable cookie processing
If set to True, cookie data should be included in session segment analysis.

Convert dates and times to defined TeaLeaf timezone
Determines the time zone in which the analysis data is stored.

Maximum number of permitted search server connections
The maximum number of connections the RSE service opens to Search Server at any time.

- Don't change this setting unless directed by Tealeaf.

Maximum number of sessions that can be processed simultaneously
The number of sessions that are processed in parallel.

Maximum number of sessions queued in memory for processing by a job
The size of the in-memory session queue.

- Don't change this setting unless directed by Tealeaf.

Number of unwritten data rows to accumulate before writing
The size that memory tables should reach before being written to SQL Server.

- Don't change this setting unless directed by Tealeaf.

The From address in emails from the Resultset Extractor
The From address used for emails sent by the RSE service.

Process URL fields
Set to True if URL Field information should be included in session segment analysis.

Enable RSE Service to create default definitions for unknown Search Servers
Set to True if the RSE service should create a default definition and attempt to contact Search Servers about which it does not know.

Database connection timeout in seconds
Defaults to 60 seconds.

Number of days to keep Resultset history
The number of days to store results sets.

Time in seconds before checking for Resultset history to remove
The interval in seconds between checks for resultset history to remove.

Field Settings

The RSE Field Settings enable you to pull specific values out of the sessions in a session segment by defining the field containing the desired value.

Note: This feature is for advanced users. To use this feature, the user must be familiar with the Tealeaf XML session definition.

When a session is analyzed, the RS Extractor reviews these field settings and extracts the specified values of the enabled fields.

- These values are visible on the Session Segment Analysis page on the Top Hit Field Values report in the Form Fields section. See "Analyzing Session Segments" in the *IBM Tealeaf cxResults User Manual*.

To edit one of the listed fields, click the entry. Then, click **Edit**. The following properties are available.

Property
Description

Section
The part of session XML in which the field appears.

Name The XML name of the field.

Value When `Enabled`, the field and its value are retrieved from sessions and inserted into the Top Hit Field Values report.

Aliases

This section lets you add Tealeaf Search Servers that were not defined at setup.

Property
Description

Host The hostname of the Search Server.

Port The port number on which to contact the Search Server.

Report Configuration

Use the links below to access the documentation on how to configure and schedule different types of Tealeaf reports. Scheduled reports are executed at the specified time and delivered via email in PDF or Excel format to a designated list of addresses.

Table 26. Report Configuration

Report Type	Configure	Schedule
Custom Report	"Tealeaf Report Builder" in the *IBM Tealeaf Reporting Guide*	"Report Manager" in the *IBM Tealeaf Reporting Guide*
Top Movers Report	"Analyzing Top Movers" in the *IBM Tealeaf Reporting Guide*	"Report Manager" in the *IBM Tealeaf Reporting Guide*
Scorecards (IBM Tealeaf cxView)	"Configuring Scorecards" in the *IBM Tealeaf cxView User Manual*	"Configuring Scorecards" in the *IBM Tealeaf cxView User Manual*
Dashboards (IBM Tealeaf cxView)	"Configuring Dashboards" in the *IBM Tealeaf cxView User Manual*	"Configuring Dashboards" in the *IBM Tealeaf cxView User Manual*

Note: If the number of days to retain hourly event reporting data is changed, the reports might not display current data until the system normalizes.

Portal Announcements

Portal announcements are messages displayed to all Portal users upon login. These messages can be used to display current system status, scheduled maintenance windows, or other Tealeaf-related issues.

- You can view current announcements by selecting **Help** > **Portal Announcements** in the Portal.

Enabling Announcements

To enable the display of Portal announcements, please review and configure the following IBM Tealeaf cxImpact settings

Property
Description

`Display Splash Screen on User Login`
When enabled, users see the Portal Announcements window when they first login.

- The default value is `Enabled`. See "CX Settings" on page 192.

`System Authentication Verification Announcement`
When enabled, the Portal reports an announcement message upon login if the Portal is mis-configured.

- The default value is `Enabled`. See "CX Settings" on page 192.

See "CX Settings" on page 192.

Configure Announcements

To configure:

1. Login to the Tealeaf Portal as an administrator.
2. From the Portal menu, select **Tealeaf** > **Portal Management**.
3. Click the Announcements link in the left navigation panel.
4. To display a list of all currently configured announcements, select **Configure**.

Currently active announcements are marked with a blue dot.

New / Edit

How to create and edit an announcement.

To create or change an announcement, click **New**, or select an existing announcement and click **Edit**.

Setting Description

Message
Up to 256 characters.

- You can force line breaks in your announcement by adding `\n` to the end of each line. For example:

```
This is line 1. \n
This is line 2. \n
Line 3. \n
```

- No other form of formatting is supported.

Starts This value is the time at which the portal begins to display the announcement to users.

- `Date` - Click this field to open a calendar.
- `Time` - Click this field to open a time selector.

Expires
This value is the time at which the Portal stops displaying the announcement to users.

- `Date` - Click this field to open a calendar.
- `Time` - Click this field to open a time selector.

Up / Down

To change the order in which announcements are displayed to users, select an announcement and click **Up** or **Down**.

Refresh

Click **Refresh** to update the list and order of announcements from the Portal.

Portal Logs

In the Portal Management page, the Logs section displays the different types of logs created by both the Portal and the Data Service. Each section has its own layout and UI.

- Logs that are displayed in the Portal are extracted from and maintained in the database.

- Logs available for download are typically stored as log files on the selected Tealeaf server.

When downloading log files larger than 5MB, the following error may be displayed:

```
Data file exceeds the maximum file size.
```

These log files must be acquired directly the appropriate directory of the server where they are maintained.

To access the Portal logs:

1. Login to the Portal as an administrator.
2. From the Portal menu, select **Tealeaf** > **Portal Management**.
3. The Portal Management page is displayed. In the left panel, click **Logs**.
4. Select the type of log to display.
5. In the right panel, you can filter the displayed log using the following drop-downs:

Drop-down	Description
Filter	Display only the log messages for the selected filter. • To display messages from all logs for the selected log type, select **<All**.
Type	Displays only the messages for the selected logging level. • To display all log messages for the selected log, select **<All**.

Global System Logging Level

The Tealeaf global logging level can be configured through TMS.

Note: By default, the global logging level is configured to report errors and status messages only. Changing this value can significantly increase or decrease the volume of log messages and the size of log files.

See "TMS WorldView Tab" on page 262.

Log File Management

By default, Portal log files are automatically rolled during the day if a log file exceeds 20MB.

- The current log filename is appended with a timestamp (in `HHMMSS` format) and saved.
- New log file data for the current date is added to the default log filename for the date. The last log data for the date is always stored in a file that does not include a timestamp.

Reporting Service, Data Collector and Query Server

These logs are maintained in the database.

These sections display recent entries created by the respective component:

- **Filter** - Select the component or service by which to filter the log

- **Type** - The type of messages to display: `Status`, `Informational`, or `Warning` or `Error`

  Note: When a Type is selected, the refreshed display includes all messages from that type only.

To refresh the display, click **Refresh**.

Data Collector log files

When any of the following display filters is selected, the listed Data Collector log is displayed.

Filter **Data Collector log**

Data Collector
`TLDataCollector*.log`

Data Collection
`TLDataCollection*.log`

Data Collection Status
`TLTDataCollectorStatus*.log`

Data Manager
`TLDataManager*.log`

Fact Aggregation
`TLFactAggregation*.log`

Performance Aggregation
`TLPerformanceAggregation*.log`

Top Movers
`TLTopMovers*.log`

Data Service

The Data Service section displays the log entries stored by the Data Service in the Tealeaf database. You can review the timestamp for when the event occurred, the type of event, and the specific log entry.

- To refresh the display, click **Refresh**.
- You can review the physical log files on the server. See "Managing Tealeaf Servers" on page 115.

RS Extractor

The RS Extractor section displays log entries from the database for the selected date created by the RS Extractor service:

- **Time**
- **Severity**
- **Component**
- **Log Entry**
- To see log entries from a different date, select the date from the Date drop-down.
- To refresh the display with the configured parameters, click **View**.

Session Segments

The Session Segments section displays a list of the currently analyzed session segments. It can be filtered by user or session segment status. Selecting a session segment displays the associated log entries.

Note: Segment analysis is a component of IBM Tealeaf cxResults. IBM Tealeaf cxResults is no longer available as a newly licensed product as of Release 8.7. Customers that licensed IBM Tealeaf cxResults in Release 8.6 and earlier may continue to use and receive support for the product in Release 8.7 and later. For more information, please contact Tealeaf Customer Support .

- To filter the list based upon the owner of the segment, select the username from the Owner drop-down.
- You may also filter the list based upon status. Make a selection from the Status drop-down to display only the segments with the selected status.
- To refresh the list, click **Refresh**.
- **ID**
- **Name**
- **Created**

For more information on analyzing session segments, see "Analyzing Session Segments" in the *IBM Tealeaf cxResults User Manual*.

Portal History

The Portal History section displays logs for the Portal activity for the current user for the current session.

For the Database History views, a list of the SQL calls to the selected database, detailing the Method, Page, Tab, SQL, Start and End Times, Time Elapsed, Query Type, and Error (if applicable) is displayed. This list can be used to view any SQL errors from the current section that are related to the selected database. The available databases are the System, Reports, RS Extractor, and Statistics databases.

These views of the data are described below.

View **Description**

AJAX Call History
A list of the different AJAX calls made by the Portal, grouped by the page from which they were called.

System Database History
A list of the recent calls to the TL_SYSTEM database, in reverse chronological order.

Reports Database History
A list of the recent calls to the TL_REPORTS database, in reverse chronological order.

RSE Database History
A list of the recent calls to the TL_RSEXTRACTOR database, in reverse chronological order.

Statistics Database History
A list of the recent calls to the TL_STATISTICS database, in reverse chronological order.

Tealeaf Status Audit Log Events

The Tealeaf Status feature can read Audit log events from the Portal.

Selecting Tealeaf Status Audit Log Events displays the tree of Audit Events, which can be individually selected or selected as a group for Tealeaf Status reporting.

View Tealeaf Status

The View Tealeaf Status command executes the Tealeaf Status check and displays the generated report in a new window.

- See "Tealeaf Status Report."

Tealeaf Logs and Configuration Files

In this section, you can select the log files to download into `.zip` format for a predefined preceding number of days.

- Downloaded files may be in text or XML format.

1. Select the Tealeaf components whose log files you wish to download.
2. Select the time period of log data to collect.

   Note: Log files that are larger than 5MB will not be included in the download `.zip` file.
3. You can optionally add a password to the `.zip` file.
4. Click **Download**.

Tealeaf Status Report

The Tealeaf Status report generates current metrics from each Tealeaf server configured in the Portal Management page. These incremental snapshots can be captured at intervals and can be used to track system changes over time or when specific modifications are made to it.

The Tealeaf Status report is a useful mechanism for troubleshooting issues in your Tealeaf solution. Tealeaf Status can be scheduled to be run periodically or can be manually executed from the command line.

- For more information on command line execution, see "Generating TealeafStatus at the Command Line" on page 233.

  Note: The Tealeaf Status report replaces the Portal Status report, which has been deprecated. Legacy customers may still configure and execute Portal Status jobs. See .

Enabling Tealeaf Status Report

The Tealeaf Status report is enabled and configured through the Tealeaf Scheduling Service. See "Configuring the Scheduling Service" in the *IBM Tealeaf CX Configuration Manual*.

Configuring Tealeaf Status Report

This section describes how to configure a status report for your IBM Tealeaf servers.

Queried Servers

After the report is enabled, report data is automatically gathered for all Tealeaf servers that are identified in the Portal. If the list of active servers in the Portal Management page reflects the sources for the Tealeaf Status report, no further configuration is required.

Note: Tealeaf Status report collects status data only from the servers that are listed and active in the Portal Management page. You must create entries in the Portal Management page for each server from which you wish Tealeaf Status report to collect status information.

To review the currently active Tealeaf servers:

1. Login as a Tealeaf administrator.
2. Select **Tealeaf** > **Portal Management**.
3. Click **Tealeaf Servers**.
4. Click the Manage Servers link.
5. The list of Tealeaf servers is displayed. When you generate a Tealeaf Status report, data is queried from all servers listed as Active.
6. To create a reference to a Tealeaf server that is not on the list, click **New**.
   - See "Managing Tealeaf Servers" on page 115.

Adding Optional Servers

In some situations, a server may be configured as a failover option if another server goes down. The Tealeaf Status report can be configured to look for these additional backup servers.

You may use the following steps to add optional servers to the TealeafStatus configuration file.

1. Login as a Tealeaf administrator.
2. Select **Tealeaf** > **TMS**.
3. Click the WorldView tab.
4. Select the Scheduling Service node.
5. Click **TealeafStatus configuration file**.
6. In the Config Actions panel, click **View/Edit**.
7. At the bottom of the file, you should see the following configuration information.

   Note: If this configuration information is not available and you are using a build that supports this optional configuration, you may copy and paste the following into your configuration file.

```
#
# Server Types:
#
#Processor
#DataCollector
#DataServer
#cxConnect
#Tms
#ReplayServer
#ReportService
#Rse
#SelectiveArchive
#Hbr
#Database
#Pca
```

```
#Verify
#Visitor
#Report
#Tli

#
# AdditionalServers:
    Server_<n>=[ServerName],[ssPort],[ServerType],[ServerType],...
#
#Note: Can be multiple types.
#
# Where:
#   <n>                 = a number, each one must be unique.
#   [ServerName]        = Name of the server
#   [Title]             = Server Title for TealeafStatus
#   [Port]              = Port number of server
#   [,ServerType]       = Server Type ie: DataCollector
#   [,ServerType]       = Server Type ie: cxConnect
#
 [AdditionalServers]
#Server_1=ServerName,DataCollector(warm),19000,DataCollector
```

8. Under the [AdditionalServers] section, you can add new servers by doing the following:
   a. Copy the Server_1 line.
   b. Paste it as a new line.
   c. Remove the hash mark (#) at the front of the line.
   d. Change the settings for the new server as needed.

      Note: An optional server can be marked as a server of multiple types. Add new server entries as comma-separated text values from the Server Types list at the top of the text above.
9. When you are satisfied with your changes, click **Save**.
10. Submit the job to TMS.
11. When Tealeaf Status runs, these additional Tealeaf servers are added to the list of servers Tealeaf Status queries.

Note: The ServerName reference does not need to be listed as a server in the Portal Management page. If it does not exist there, you do not need to add it. These additional servers are used by the Tealeaf Status report only.

Generating Tealeaf Status Report

1. Login as a Tealeaf administrator.
2. Select **Tealeaf** > **Portal Management**.
3. Click **Logs** in the left navigation pane.
4. Click the View Tealeaf Status link.

Generating TealeafStatus at the Command Line

You can execute Tealeaf Status reports at the command line when needed.

The file to execute is the following:

```
<Tealeaf_install_directory>\Tools\TealeafStatus.exe
```

The command line accepts the following parameters.

Parameter
Description

-? Get help from `TealeafStatus.exe`.

-dataserver
Server host of the Tealeaf Data Service

-dataserverport
Port number through which to communicate with the Tealeaf Data Service on the above host

-general
Generate a general Tealeaf Status report

-verbose
Generate a verbose Tealeaf Status report

-file The filename of the outputted report

Sections of Tealeaf Status Report

This section provides an overview of the various sections of the Tealeaf Status report.

At the top of the report, you can review the physical server from which the report data was take, as well as a timestamp and Tealeaf build number associated with the Tealeaf Status report.

Tealeaf Server Status

This section identifies the current status of the Tealeaf servers on the physical server that was queried.

Field **Description**

Server The server name for the Tealeaf server.
- A value of `localhost` indicates that the server is stored on the Portal Server.

Type The type of Tealeaf server.

Status Short status message.
- `OK` denotes that the server is functioning.

Tealeaf's System Throughput

This section indicates the current flow rate of Tealeaf data through the server.

Flow rate

General statistics on data processing through the Tealeaf system.

- This data is gathered from all active Canisters (Processing Servers) in the Tealeaf system.

Field **Description**

Sessions per hour
Number of sessions processed in the last hour

Current Hits per second
Number of hits processed per second

Current Bytes per second
Number of bytes processed per second

Sessions last 24 hours
Number of sessions processed in the last 24 hours

Max Hits per second
Maximum number of hits processed per second

Max Bytes per second
Maximum number of bytes processed per second

Decoupler Overview

Overview statistics on the Decouple session agent on the server. This session agent manages increases in traffic volume, including the spooling of hits for processing at a later time.

- See "Extended Decoupler Session Agent" in the *IBM Tealeaf CX Configuration Manual*.
- Some of these statistics are available in the All Servers report for the DecoupleEx status report. See “System Status” on page 8.

Field **Description**

Server The instance of the Tealeaf pipeline monitored by the session agent

Can Check
When ON, the Canister checking monitoring process is enabled.
- This setting can be enabled through the session agent configuration. See "Extended Decoupler Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

Status Current status message of the session agent. Values:
- Memory - hits are being processed in memory only.
- Disk - hits are being spooled to disk because memory is full.
- UNKNOWN - the status of the DecoupleEx session agent is not known. The session agent is not operational.

Hits In
Rate per second of hits reaching the session agent

Can Hits In
Rate per second of hits reaching the canister

Hits in Memory
Number of hits in memory at time of report execution

Hits on Disk
Number of hits on disk at time of report execution

Time on Disk
Time in seconds that hits on disk have spent in that location

Canister Overview

This section contains canister status information for the indicated server. Key statistics in this section are the number of Unindexed sessions and unevaluated hits.

- See “System Status” on page 8.

Field **Description**

Server The name of the Processing Server

CPU Usage %
Current CPU usage as a percentage of capacity

Memory The available memory on the server

Memory Usage
The current memory in usage on the server

Active Sessions
Number of active sessions stored in memory

UnIndexed Sessions
Number of completed sessions waiting for the indexing process

UnEval Hits
Number of hits in the canister that have not yet been evaluated

Tealeaf CX Usage

This section publishes the currently selected audit log events for Tealeaf Status.

• See "Portal Logs" on page 227.

Field Description

Description
Display name for the audit log event

Total Total count of events in last 24 hours

Succeeded
Count of successful events in last 24 hours

Failed Count of failed events in last 24 hours

Unique Users
Count of unique users who triggered the event

Tealeaf Server: Name

For each physical server, a separate section in the Tealeaf Status report provides details on the individual Tealeaf servers installed on it and other aspects of system performance.

Field Description

Tealeaf server type
The types of Tealeaf servers stored on the physical server

64 Bit OS
When `true`, the server is using a 64-bit version of the operating system

Virtual OS
When `true`, the server is using a virtual operating system.

Note: For optimal performance, Tealeaf software should be installed on dedicated hardware. Tealeaf does provide limited support for virtual machine (VM) environments. See "CX Pre-Installation Checklist" in the *IBM Tealeaf CX Installation Manual*.

Canister Status:

For each server hosting Tealeaf canisters, this section provides a status summary of the canister performance at the time of report execution.

• See "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

General

Field Description

Server Time
The local time on the canister at the time of report execution

Last CanTrim
The timestamp when the CaniTrim operation was last performed, if ever.

- The CanTrim process removes aged sessions from the Long-Term Canister. CanTrim is enabled through TMS. See "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

Time Zone
The time zone where the canister is located

Canister Database Version
The Tealeaf build number of the ctree database in use on the Canister server.

Sessions

The Sessions section provides statistics on session counts passing through the Canister.

Field Description

Total Sessions (STC)
Total sessions in the Short-Term Canister (active sessions) at time of report execution

Total Sessions (LTA)
Total sessions in the Long Term Archive (completed sessions) at the time of report execution

Sessions Waiting for Evaluation
This metric is deprecated.

Sessions Waiting for Timeout
This metric is deprecated.

Sessions Waiting to be Archived
Number of completed sessions waiting to be written to disk.

Sessions Not Indexed
Number of completed sessions that have yet to be indexed

- See "Configuring CX Indexing" in the *IBM Tealeaf CX Configuration Manual*.

Hits/Misc

This section includes statistics on hits and CPU utilization.

Field Description

Evaluated Hits
Number of evaluated hits in the STC

Non-Evaluated Hits
Number of non-evaluated hits in the STC

CPU Utilization
Current rate of CPU utilization at time of report execution

CPU Utilization (Last 10 Minutes)
Rate of CPU utilization over the previous 10 minutes

CPU Utilization (Last Hour)
Rate of CPU utilization over the previous hour

Last Trim Canister time
The timestamp for when the last CanTrim operation has been performed

Current Facts
Average number of facts processed

Current Aggregates
Average number of aggregations

Current Aggregate Keys
Average number of aggregation keys

Disk (Canister)

In this section, you can review disk storage data for the canisters.

Field Description

Canister Size on Disk
Current size of the LTC on disk

Disk Size
Storage capacity of the disk where the LTC is stored

Free Disk Size
Available disk space on the disk

Canister Path
Path to the session data files on the canister.

- Default is `<Tealeaf_install_directory>\Canister\Canister.dbs`

Index Path
Path to the index data files on the canister.

- Default is `<Tealeaf_install_directory>\Canister\Indexes`

Disk (Index)

In this section, you can review disk storage data for the indexes.

Field Description

Index Size on Disk
Current size of the index files on disk

Name (LSSN file)

This section provides status information for the canister files stored on the selected server.

- Session data files are stored in the LSSN format.

Field Description

Name Filename of the LSSN file

Sessions
Number of sessions contained in the LSSN file

Size Size of the LSSN file

Decouple Status:

The Decouple section contains a subset of statistics reported in the DecoupleEx status report for individual servers.

- See "System Status" on page 8.

TMS Status:

This section contains a list of configuration changes executed through TMS over the past twenty-four hours.

- The name of the TMS master is listed in the section title bar.

Update configurations

This table identifies all of the files controlled by TMS that have been modified over the past 24 hours.

Field Description

Filename
The name of the file that was changed

Component
The TMS component that was changed

Modified
The timestamp for when the changed was committed to TMS

Version
The version number of the committed change

External configurations

A list of filenames that were changed on servers other than the TMS server via TMS jobs.

TMS jobs

This table indicates the TMS jobs that have been started over the past 24 hours.

Note: After 24 hours has elapsed from the time when the job was created, the job is removed from the Tealeaf Status report.

Field Description

Job Type
The type of job executed.

ID The TMS internal identifier for the job.

Description
The user-friendly description for the job.

Status An indicator of whether the job succeeded or failed. Possible values:

- 1 - Queued
- 2 - Scheduled
- 3 - Waiting on Dependency
- 4 - Processing
- 5 - Rescheduled

- 6 - Finished
- 7 - Finished (with errors)
- 8 - Failed
- 9 - Interrupted
- 10 - Cancelled

Scheduled
The time at which the job was scheduled to execute.

Services:

On the current server, this section lists the Tealeaf services that are currently available and their current operating status.

- These services can be enabled and disabled through TMS. See Chapter 4, "Tealeaf Management System," on page 259.

Tealeaf CX Usage - last 24 hours:

For the preceding 24 hours, you can review the number of times common user activities were executed by Tealeaf users.

- For a list of the audited user events, see "Portal Logs" on page 227

Application Events:

This section contains a summary of all application log events logged for the Tealeaf system.

Depending on the logging level and volume of log entries, this section may be an easy means of reviewing the Tealeaf application logs for the selected server.

- See "Portal Logs" on page 227.

Field	Description
Type	Type of application event message
Time	Timestamp of event message
Source	Tealeaf component that generated the message
Category	Category identifier for the message
ID	Message identifier for the message
Message	Full text of the message

System Events:

This section contains a summary of all system log events logged by the hosting server.

Field	Description
Type	Type of system event message
Time	Timestamp of event message
Source	Windows component that generated the message

Category
Category identifier for the message

ID Message identifier for the message

Message
Full text of the message

Data Aggregation and Retention

The Tealeaf reporting system is driven by data collected from the Tealeaf Processing Servers. Collected data consists of statistical information that is generated while sessions are processed. The data is generated in one-minute buckets and extracted into the reporting database every five minutes.

- The actual text of the sessions remains on the Processing Server and is not migrated to the reporting databases.

The collected data is aggregated into two types of reporting data:

- hourly
- daily

After data has aged a pre-defined period of time, the data is removed from the database so that it can be kept to a manageable size.

- For more information on retaining data beyond the retention period, see "Retaining Data after Database Expiration" on page 248.

The length of the retention period is positively correlated to the size of the database; retaining more data can result in a very large database, particularly in the tables that store hourly data.

This section provides some guidelines in configuring data retention and aggregation.

Session Data Retention

Session data is stored on the local disk for each Processing Server (Canister) in your environment. It is not retained in a Tealeaf database.

Configuration of session data retention is managed for each Canister through the Tealeaf Management System.

- See "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.
- See Chapter 4, "Tealeaf Management System," on page 259.

Data Collector Service

The Data Collector Service manages the collection, aggregation, and trimming of the base reporting data, as well as the calculation of any predefined reports. It is enabled and disabled through TMS.

- Data collection and aggregation processes run independently of each other.

  Note: If the Data Collector Service is disabled, Portal reports are not updated.
- See Chapter 4, "Tealeaf Management System," on page 259.

Data Collection

Data Collection is enabled and disabled through the `Data Collection` setting.

Note: This setting should not be disabled during normal operations. It is available primarily for debugging purposes.

Setting Description

`Data Collector Logging Level`
Specify the logging level for the Data Collector only.

- Don't change this setting unless directed by Tealeaf. Levels:
- 0 - `none`
- 1 - `Error` (default value)
- 2 - `Warning`
- 3 - `Info`
- 4 - `Detail`
- 5 - `Status`
- 6 - `Trace`
- 7 - `All`

Note: Status level messages always appear in the log for any non-zero logging level.

This value overrides the system logging level, which can be configured through TMS. See "TMS WorldView Tab" on page 262.

- See "CX Settings" on page 192.

Data Aggregation

This section describes performance, logging, and managing the data aggregation.

cxImpact

Aggregation of IBM Tealeaf cxImpact data for reporting purposes is managed through the following settings.

Setting Description

`Data Aggregation`
Set this value to `Enabled` to enable data aggregation of collected data for reporting purposes.

Note: This setting should not be disabled during normal operations. It is available primarily for debugging purposes.

- Don't change this setting unless directed by Tealeaf.

`Data Aggregation - Max Concurrent Threads`
The maximum number of concurrent threads that can be used for aggregating base reporting data.

See "CX Settings" on page 192.

Performance Data

Performance data includes time values on the visitor's client browser, the server response times, and the connection times for the visitor. These data are aggregated and reported in the Performance reports.

- See "Analyzing Performance" in the *IBM Tealeaf Reporting Guide*.

Setting Description

`Data Aggregation - Performance Data`
When `Enabled`, data on client performance, response times, and connection times is aggregated for reporting purposes. See "Analyzing Performance" in the *IBM Tealeaf Reporting Guide*.

`Data Aggregation - Max Concurrent Performance Threads`
The maximum number of concurrent threads that can be used for aggregating performance data.

See “CX Settings” on page 192.

Dimension Logs

For any dimension flagged for logging, the Canister gathers all observed values for the dimension from the capture stream. Those values are written to a database table once per hour.

By default, the Data Collector polls the Canister ten minutes past the hour to gather new values from the Canister dimension log and inserts newly detected ones into the database. These values can then be used in specifying whitelists, blacklists, and group lists for dimensions that reference them.

The interval and time when the Data Collector values are aggregated can be configured as needed. The following configuration options are available:

- Enable/disable aggregation of values
  - You should not need to disable aggregation of values, unless you are attempting to troubleshoot an issue with the Canister or Data Collector.
- Interval of Data Collector aggregation (`Dimension Log Aggregation Interval`):
  1. Hourly - When log values are aggregated on an hourly basis, the Data Collector process runs at ten minutes after the hour (1:10, 2:10, etc.).
  2. Daily or Weekly - When log values are aggregated on a daily basis, you may configure the time of day when the log values are collected (`Dimension Log Aggregation Time`).

     Note: When dimension log aggregation is scheduled for a daily or weekly basis, the number of log values that must be collected and aggregated may impact system performance. Tealeaf recommends setting this value to an off-peak hour. By default, it is configured to execute at `3:00` in the morning.

Note: The logging of dimension values must be enabled for each dimension that you create. After these values have been inserted into the database, they can be added to the whitelists, blacklists, or group lists for the referencing dimension through the Tealeaf Event Manager. See "TEM Dimensions Tab" in the *IBM Tealeaf Event Manager Manual*.

The following settings are used to configure the Data Collector settings for dimension log values.

Setting Description

Dimension Log Aggregation
When Enabled, Tealeaf dimension values are aggregated from log entries at predefined intervals.

Dimension Log Aggregation Interval
When Dimension Log Aggregation is enabled, this setting defines the time interval between checks of the logs for reference values.

Dimension Log Aggregation Time
When Dimension Log Aggregation is enabled and Dimension Log Aggregation Interval is set to Daily, this setting defines the 24-hour time when the review of the logs is executed.

See "CX Settings" on page 192.

Managing Data Aggregation Performance

Depending on the volume of data retained in your reporting database, the periodic data aggregation runs performed by the Data Collector can impact performance and consume unnecessary space. Performance in the Report Builder can also be impacted.

To improve performance, Tealeaf provides the following controls to define the frequency and time of day that the hourly data is aggregated to the daily level.

Setting Description

Data Aggregation - Daily Data Processing
Describes the available settings for the date range over which daily data aggregations occur. Options:

- Hourly through Current Hour - Aggregate once per hour, in first data collector run of the hour, using most current available hourly data
- Daily Through Previous Day - Aggregate daily data once per day, using hourly data up through the previous day
- Daily Through Start of Hourly Retention Period - Aggregate daily data once per day using hourly data up to the beginning of the hourly data retention period

  Note: The default value, Daily Through Start of Hourly Retention Period, optimizes the aggregation process for performance of data collection. Configuring this setting to a value other than Hourly Retention Period retains overlapping aggregated data in the database and may impact system performance during data aggregation.

Data Aggregation - Daily Data Time of Day
This setting defines the hour of the day when the daily data aggregation run is performed, if the daily data aggregation is set to occur on a daily basis. By default, it is configured to run at 2:00 in the Tealeaf system timezone.

Note: Tealeaf recommends configuring the daily data aggregation run to be performed during an off-peak hour.

See "CX Settings" on page 192.

Data Trimming

Data removal can be configured through the use of the following settings:

Setting Description

Data Trimming - Canister Data
When Enabled, the un-aggregated event data is trimmed based on the configured Data Collector settings.

Note: When this value is set to Disabled, no trimming occurs at all. The SQL Server database can grow without limit.

Data Trimming - Canister Data Immediate Trim
When Data Trimming - Canister Data is Enabled, enabling this parameter forces the canister data tables to be immediately trimmed after the Data Collector has aggregated the data in them.

Note: If errors are encountered during an aggregation operation, the data in with which the errors are retained. Re-aggregation and trimming is attempted during normal canister trim operations.

- If Data Trimming - Canister Data is disabled, this parameter is ignored.
- Don't change this setting unless directed by Tealeaf.

Data Trimming - Interval
Determines the interval at which the Data Service trims data from the database 0 - None, 1 - Hourly, 2 - Daily, 3 - Weekly.

Data Trimming - Max Batch Size
The maximum number of records trimmed from the reporting or canister data tables in any single delete statement executed as part of a trimming operation.

Data Trimming - Reporting Data
When Enabled, the reporting data is trimmed based on the configured Data Collector settings.

Note: When this value is set to Disabled, no trimming occurs at all. The Reports database can grow without limit.

Data Trimming - Statistics
When Enabled, the Tealeaf statistics data is trimmed based on the configured Data Collector settings.

Note: When this value is set to Disabled, no trimming occurs at all. The Statistics database can grow without limit.

Data Trimming - System
When Enabled, the user activity logs in the TL_SYSTEM database are trimmed based on the configured Data Collector settings.

Note: When this value is set to Disabled, no trimming occurs at all. The related System database tables can grow without limit.

Data Trimming - Time for Daily Trim
If the value for Data Trimming - Interval is Daily, then this value defines the time of day at which the reporting data is trimmed. Time is based on the Tealeaf system time zone. It should be configured for an off-peak hour.

Data Trimming - Top Movers
When Enabled, the Top Mover data is trimmed based on the configured Data Collector settings.

Note: When this value is set to Disabled, no trimming occurs at all. Top Movers data in the database can grow without limit.

See "CX Settings" on page 192.

System Data

The Tealeaf system is configured and runs off of data stored in the system database TL_SYSTEM. The data contains information on users and application and report configuration.

Reporting Data

The Tealeaf reporting system is driven by event-based data. All reporting data is based on the defined events and dimensions and is stored in the reporting database: TL_REPORTS.

The data is divided into two types.

1. **Non-dimensional event data.** This data consists of statistics and values gathered for events without respect to any dimension data.
2. **Dimension-based event data.** This data is based on the defined report groups assigned to an event. Each report group may contain multiple dimensions along with the timestamp and corresponding event counts and values.

Tealeaf System Statistics Data

In the TL_STATISTICS database, statistics are aggregated by time and Tealeaf component.

Dimension Data Trimming

To limit unbounded growth of dimension data, the Data Collector performs automated trimming of dimension values, based on Portal settings.

- See "Data Management for Dimensions" in the *IBM Tealeaf Event Manager Manual*.

Data Retention

Through the Portal, you can define the number of days that data is retained before it is purged. At install time, these settings are configured to meet the requirements of your Tealeaf solution.

Over time, the volume of web traffic and enterprise requirements may require more or less data retention in the database. Below, you can review the generalized effects of increasing the length of time that data is retained.

- Additional information on the effects of changing data retention for each type of retained data is listed in the following sections.
- For more information on retaining data beyond the retention period, see "Retaining Data after Database Expiration" on page 248.

General effects of increasing data retention

1. The size of the database grows. Depending on the type of data, this growth can be substantial, particularly for hourly data.
2. The increased volume of data negatively impacts the speed of data aggregation for the Data Collector. The larger indexes on the tables result in increased lookup and update times to the data tables.
3. The performance of Portal reporting is negatively impacted for the same reasons that data aggregation is affected. Larger indexes cause longer lookups and query processing.
4. Depending on the amount of RAM available on the server hosting the database, the increase in data may cause an increase in SQL Server paging. If the amount of data being accessed consumes all available RAM, SQL Server pages the data to disk, which results in slow query processing.

Decreasing the length of data retention results in the inverse of the above impacts. In some cases, reducing these values can significantly increase performance, especially if there are hardware or storage limitations.

Reporting data retention

The reporting data consists of data aggregated by time across all events and report groups. The size of the data depends on the number of events, their firing frequency, and the variation in values for and dimensions used in report groups.

For more complex and high volume systems, these stored data types can rapidly lead to a large volume of data. Properly setting the data retention settings is the best way to control the size of the Tealeaf databases and ensure proper system performance.

Data retention is controlled by the settings found in the Data Retention section of the IBM Tealeaf CX Settings page.

For information on the individual settings, see “CX Settings” on page 192.

System Processing Stats Data Retention

Individual components in the Tealeaf system have access to a statistical logging function, which captures stats data as hits from the Tealeaf pipeline for collection by the Data Collector. These statistical hits are then written directly to the database.

For information on the individual settings, see “CX Settings” on page 192.

User Activity Log Data Retention

All user activities executed in the Tealeaf Portal are captured as log data and stored in the Tealeaf database.

Setting Description

`User Activity Log Data-Days Retain`
Number of days to keep the Portal user activity history data.

See “CX Settings” on page 192.

Data Retention for cxResults

IBM Tealeaf cxResults provides access to a rich database repository of information on visitors to your web site.

You can configure data retention settings for both IBM Tealeaf cxResults databases to control the database size requirements of IBM Tealeaf cxResults.

Note: IBM Tealeaf cxResults is no longer available as a newly licensed product as of Release 8.7. Customers that licensed IBM Tealeaf cxResults in Release 8.6 and earlier may continue to use and receive support for the product in Release 8.7 and later. For more information, please contact Tealeaf Customer Support .

Data retention settings for the reporting and staging databases for IBM Tealeaf cxResults are governed by the following parameters in IBM Tealeaf cxResults Settings under Visitor Extract.

Setting Description

`Reporting Visitor Data-Days Retain`
: This value is the number of days of data retained for reporting purposes.

See "cxResults Settings" in the *IBM Tealeaf cxImpact Administration Manual*.

Retaining Data after Database Expiration

After data has reached its expiration date, the Data Collector purges the data from the database. It is no longer available.

To retain this data, you have the following options:

1. Integrate the Tealeaf databases with another enterprise data system. Using IBM Tealeaf cxConnect for Data Analysis, you can extract data from the Tealeaf database and insert it into another enterprise system. See "cxConnect for Data Analysis Administration Manual" in the *IBM Tealeaf cxConnect for Data Analysis Administration Manual*.

   Note: IBM Tealeaf cxConnect for Data Analysis is a separately licensed component of the IBM Tealeaf CX system. please contact your IBM Tealeaf representative.
2. Forward to the enterprise data warehouse. Using IBM Tealeaf cxVerify, you can extract session data and store them in the enterprise data warehouse for archival purposes. See "cxVerify Administration Manual" in the *IBM Tealeaf cxVerify Administration Manual*.

   Note: IBM Tealeaf cxVerify is a separately licensed component of the IBM Tealeaf CX system. IBM Tealeaf cxVerify is no longer available as a newly licensed product as of Release 8.7. Customers that licensed IBM Tealeaf cxVerify in Release 8.6 and earlier may continue to use and receive support for the product in Release 8.7 and later. For more information, please contact Tealeaf Customer Support .
3. Increase data retention settings and corresponding storage. The restrictions and costs identified in this page apply.
4. Build a custom solution.

Configuring Locale Settings

This section provides information on how to configure the language to display in the Portal and how Tealeaf processes and displays data that may vary depending on the selected language.

Configuring Language for the Portal

The Portal can be configured to display text for any of the supported languages. Language selection is primarily based on user and group configuration options.

Configuring regional and language settings

If you have installed the operating system using a non-US locale and are experiencing issues related to your locale, you should apply the Microsoft solution to apply regional and language settings to the Windows reserved accounts.

Note: The following solution from Microsoft should be applied only if you are experiencing issues related to your deployment in a non-US locale. For more information from Microsoft, see http://windows.microsoft.com/en-US/windows-vista/Apply-regional-and-language-settings-to-reserved-accounts.

Portal Login Page

Before a user has logged in to the Portal, any user-specific language settings are unknown. If needed, you can configure the language to display in the Portal login page, so that all Portal users receive the same messaging in the same language.

1. Login to the Portal as an administrator.
2. From the Portal menu, select **Tealeaf** > **Portal Management**. The Portal Management page is displayed.
3. From the left navigation panel, select the IBM Tealeaf CX Settings category.
4. Click the Users link.
5. Select the `Login Page Language (Global)` setting.
6. Scroll down. From the drop-down, select the language to display in the Portal login page.
7. Click **Save**.
8. All subsequent visitors to the login page see text displayed in the selected language.
   - See “CX User Administration” on page 165 in the “CX User Administration” on page 165 section.

DNS Computer Names

The Portal enforces proper specification of DNS names. DNS computer names can contain alphabetical characters (`A-Z`), numeric characters (`0-9`), the minus sign (`-`), and the period (`.`) only. Period characters are allowed only when they are used to delimit the components of domain style names.

- For more information, see `http://technet.microsoft.com/en-us/library/cc785282%28WS.10%29.aspx`.

Configuring Languages for User Accounts

Inside the Portal, the displayed language is determined by user and group configuration. Users inherit the language setting from their primary user group, although individual user accounts can be configured to override that selection.

By user group

Please complete the following steps to configure the language to display to members of a user group.

In the example below, the group that is configured is the `Portal User` group, which is the first group to which all users are assigned. By default, this user group becomes the primary user group. Although it can be overridden for individual users, you should configure the language associated with the `Portal User` group to be the one in most common use among your Tealeaf user community.

1. Login to the Portal as an administrator.
2. From the Portal menu, select **Tealeaf** > **Portal Management**. The Portal Management page is displayed.
3. From the left navigation panel, select the IBM Tealeaf CX User Administration category.
4. Click the Groups link.
5. Select the `Portal User` group.
6. From the Default Language drop-down, select the language to display to members of the `Portal User` group after they have logged in to the Portal.
7. You may also wish to select a different Default Date Format.
8. Click **Save**.
9. All Tealeaf users who are members of the `Portal User` group see text inside the Portal displayed in the selected language and date format.
   - See “CX User Administration” on page 165 in the “CX User Administration” on page 165 section.
10. The above steps should be applied to other IBM Tealeaf cxImpact user groups.
11. If you have licensed IBM Tealeaf cxView or IBM Tealeaf cxReveal, you may also need to apply similar steps to the above to change the language and date format for primary user groups for those products.
    - Repeat the above steps, selecting the IBM Tealeaf cxView User Administration link or IBM Tealeaf cxReveal User Administration link in the Portal Management page as needed.
    - See “cxView User Administration” on page 184 in the “cxView User Administration” on page 184 section.
    - See "cxReveal User Administration" in the *IBM Tealeaf cxReveal Administration Manual* section.

By individual user

By default, individual users inherit the language settings from their primary group default settings.

- For example, for most IBM Tealeaf cxImpact users, their primary user group is configured to be `Portal User`. If you performed the above configuration for that user group, these steps may be necessary only to manage exceptions.

Please complete the following steps if you need to override the language settings configured for an individual user's primary default group.

1. Login to the Portal as an administrator.
2. From the Portal menu, select **Tealeaf** > **Portal Management**. The Portal Management page is displayed.
3. From the left navigation panel, select the category in which the user is located:
   - `CX User Administration`
   - `cxView User Administration`

- `cxReveal User Administration`

4. Click the Users link.
5. Select the user to modify.
6. For the Primary User Group setting, verify that the user has been properly assigned to the appropriate primary user group.
   - An improperly specified primary user group may result in the wrong language displayed to the user.
7. If the Primary User Group setting is correct, select the user's language from the Language drop-down.
8. You may also configure date format options as needed.
9. If you have made changes, click **Save**.
10. When the user next logs in, the changes are applied.
    - See “CX User Administration” on page 165 in the “CX User Administration” on page 165 section.
    - See “cxView User Administration” on page 184 in the “cxView User Administration” on page 184 section.
    - See "cxReveal User Administration" in the *IBM Tealeaf cxReveal Administration Manual* section.

My Settings

Individual users may also configure language and date format selections for themselves.

- See "My Settings" in the *IBM Tealeaf cxImpact User Manual*.

Configuring Time Zone

This section describes how time zones are used and configured within IBM Tealeaf.

Configuring Tealeaf Time Zone

Tealeaf utilizes a single time zone as the basis for timestamping. The system time zone is available through the Tealeaf Management System.

Note: You should configure the Tealeaf time zone as soon as you have completed an installation or an upgrade to Release 8.x or later. See "Configuring the System Timezone" in the *IBM Tealeaf CX Configuration Manual*.

For more information on TMS, see Chapter 4, “Tealeaf Management System,” on page 259.

Time Zones for Search

For each user group, you may configure a time zone to applied when performing searches. This time zone overrides the Tealeaf time zone for purposes of resolving the day when a session occurred. For example, if you select to search for sessions that occurred "today," the returned sessions occurred since the start of midnight for your group or individual user time zone settings.

Group time zones

To change the time zone setting for a user group, you can follow the instructions for modifying language in this section. In the same location in the Portal Management page, the time zone setting for the group can be selected from the `Default Time Zone (used in Search)`.

- See “By user group” on page 250.

User time zones

To change the time zone setting for an individual user, you can follow the instructions for modifying language in this section.

In the same location in the Portal Management page, the time zone setting for the user can be selected from the `Default Time Zone (used in Search)`. This setting overrides the group time zone setting for the user's primary group.

• See “By individual user” on page 250.

User-configured time zone

Individual users may configure their own time zone setting. For more information, see "My Settings" in the *IBM Tealeaf cxImpact User Manual*.

Configuring Currency

By default, the Tealeaf Portal displays values formatted as currency in U.S. format.

For example, one hundred dollars and ninety-nine cents is represented as the following:

```
$100.99
```

Note: Currency values are not displayed in the Tealeaf Portal. The following sections enable you to define how currency values are identified, even though the identifiers may not be displayed as part of the value.

Configuring the Canisters to Recognize Currency Identifiers

As needed, you can configure the Processing Servers (Canisters) to parse numeric data as currency, using the decimal separator of your choice. Please complete the following steps.

1. Login to the Portal as an administrator.
2. From the Portal menu, select **Tealeaf** > **TMS**. The Tealeaf Management System is displayed.
   • See Chapter 4, “Tealeaf Management System,” on page 259.
3. From the View drop-down, select **Servers**.
4. Open one of the server nodes where a Canister is installed.
5. Click the Canister node.
6. Click **Canister configuration**.
7. In the Config Actions panel, click **View/Edit**.
8. Configure the following properties:
   a. `Decimal Character for Numbers`: Set to the appropriate value to signify a decimal in captured data.
   b. `Parse Numerics as Currency`: When set to `true`, numeric values in the capture stream are treated as currency values. The `Decimal Character for Numbers` is applied to the data.
   c. To save any changes, click **Save**.
9. Push the changes to all servers.

Supported Languages

Note: All languages supported for use in the Tealeaf Portal are available through the Language settings in the Portal Management page. If you do not see your preferred language, please submit a new feature request with Tealeaf http://support.tealeaf.com.

Encoding

Tealeaf supports multiple methods of encoding, which may be configured to support various languages.

- See "Internationalization Support" in the *IBM Tealeaf CX Installation Manual*.

Logging

The following logs contain information on the locale settings in use for the Portal.

Data Service log

Data Service log entries may contain the following information:

```
OS Name: Microsoft Windows Server 2008 R2 Enterprise
OS Service Pack:
OS Version: 6.1.7600.0
OS Architecture: x64
OS UI Culture: en-US - English (United States)
OS UI Culture NumberFormat: Group=',' Decimal='.'
Thread Culture: en-US - English (United States)
Thread Culture NumberFormat: Group=',' Decimal='.'
Thread UI Culture: en-US - English (United States)
Thread UI Culture NumberFormat: Group=',' Decimal='.'
```

Entry Description

`OS UI Culture`
UI culture in use by the operating system

`Thread Culture`
UI culture for parsing input data

`Thread UI Culture`
UI culture in use for displaying data

About page

The above information is also displayed in the Portal's About page.

1. Login to the Portal as an administrator.
2. From the menu, select **Help** > **About IBM Tealeaf CX Portal**.
3. Review the values in the System Property table.
   - See “Portal Application Information Report” on page 78.

Configuring Character Encoding

Tealeaf supports the use of different types of character encoding.

- See "Internationalization Support" in the *IBM Tealeaf CX Installation Manual*.

Backing up the IBM Tealeaf event model

Through the Portal, you can define tasks to automatically create backups of critical event model data. As needed, Portal backups can be restored, returning your Tealeaf solution to a previously saved state.

Overview

To access Portal backups, select **Tealeaf** > **Portal Management** from the Portal menu. In the left navigation panel, click **Backup & Restore**.

- For more information on the types of data stored in a backup, see "Archived Event Model Data."

You can perform the following types of backups:

1. **Scheduled Backups:** On a daily basis, the Portal can be configured to perform a backup of event model data. See "Configuring Backup Settings" on page 255.
2. **On-Demand Backups:** As needed, you can execute backups through the Portal at any time. "Managing Backups" on page 256.

Through the Portal Management page, you can use a backup to restore the event model to a previously saved state.

- See "Restoring Backups" on page 256.

Limitations

Note:

1. This backup feature cannot be used to backup the event data model and restore it across versions of Tealeaf. To migrate events across Tealeaf systems and versions, you must use the import/export features of the Event Manager. See "TEM Import-Export Tabs" in the *IBM Tealeaf Event Manager Manual*.
2. This backup feature does not archive any data captured or extracted from sessions. That data is stored in the Tealeaf databases, which must be backed up through the Tealeaf Database Manager or SQL Server Management Studio. For more information on performing backups through Tealeaf, see "Backing Up Tealeaf" on page 258.

Archived Event Model Data

Data configured through the Tealeaf Event Manager are archived using the event model backup, including the following:

Note: Alert definitions are not supported for backup through this method.

1. Event definitions

   Note: When events are restored, any references in reports to events that are no longer present in the system are removed from the report. In some cases, this removal may result in empty reports, scorecards, and dashboards.
2. Hit Attribute definitions
3. Dimension data, including:
   a. Dimension and report group definitions
   b. Dimension whitelists

      Note: Non-whitelisted values are not backed up.

c. Dimension group lists

Note: During restore, whitelist and group list values are restored. Any additional values already in the system are untouched.

4. Session Attribute definitions
5. Top Mover definitions

See "Tealeaf Event Manager" in the *IBM Tealeaf Event Manager Manual*.

Backup Data Files

Backup files are compressed and encoded. They cannot be read outside of the Tealeaf environment.

Configuring Backup Settings

In the Portal Management page, you can enable daily backups and configure settings for automated daily backups of event model data, which are stored on the Report Server in the designated directory.

Note: Before scheduling automated backups, you should perform a single on-demand backup first to assess the storage requirements for each backup. Then, you can determine the space available for retaining backups on the server and configure the number retained accordingly. See “Managing Backups” on page 256.

1. In the left navigation panel of the Portal Management page, click the Backup & Restore category.
2. Click **Backup Settings**.
3. The Backup settings are displayed. To change any of the following settings, click **Edit**.

Setting Description

`Backup Directory`
Directory inside the Tealeaf install directory where backups are stored.

Note: Only UNC and absolute file paths are supported.

`Backup Time`
Time of day when backups are created. By default, the time is set to 12 midnight (`00:00`).

`Number of Backups Retained`
Number of days of backup data that are retained in the `Backup Directory`.

Note:

1. This setting is not enforced until scheduled backups have been enabled. If scheduled backups are disabled, all backups in the directory are retained.
2. If you wish to retain event model data backups for longer than the number of days specified in this setting, you can download the backup from this screen to your local system for safekeeping. Older backups are purged from the directory by Tealeaf.

`Scheduled Backups`
Set this value to `Enabled` to perform scheduled backups.

Note:

1. If scheduled backups are disabled, all backups in the directory are retained. If you have more than the `Number of Backups Retained` stored in the directory, enabling this setting causes the oldest backups to be purged when the first scheduled backup is executed.
2. This setting must be enabled to perform scheduled backups of event model data.

To save your changes, click **Save**.

When backups are executed, you can review them through the Portal. In the Backup & Restore panel, click the Backup/Restore link. See "Managing Backups."

Managing Backups

Note: If backups take more than two minutes, you should consider scheduling during off-peak hours.

At any time, you can archive event model data to the specified backup directory through the Portal. For example, if you are experimenting with Portal settings or performing significant modifications to your event definitions, you may wish to perform an on-demand backup before beginning to make these changes.

- For more information on the directory where backups are stored and other settings, see "Configuring Backup Settings" on page 255.

1. From the Portal Management page, click the Backup & Restore link.
2. Click the Backup/Restore link.
3. The backup is generated and stored in the backup directory. An entry is listed in the Portal Management page.
   - On the Report Server, this backup is stored in the specified backup directory and retained until a specified number of backups have been stored there.

     Note: If you wish to retain event model data backups for longer than the number of days specified in this setting, you must configure an task external to Tealeaf to move these backups to a different directory. Older backups are purged from the directory by Tealeaf.
   - See "Configuring Backup Settings" on page 255.

Creating On-Demand Backups

To create a backup right now, please complete the following steps:

Note: If backups take more than two minutes, you should consider scheduling during off-peak hours.

1. In the Portal Management page, click the Backup & Restore category.
2. Click **Backup/Restore**.
3. Click **Create Backup**.
4. The backup is created and stored in the designated directory. It is available in the Backup/Restore list of backups.

Restoring Backups

This section describes how to restore and migrate IBM Tealeaf from a stored backup.

Backup warnings

Note:

1. Restore the system only during off-peak hours.
2. Restoring a backup automatically forces all users out of the Portal and does not save any unsaved work.

Restore steps

Note: Restoring a backup overwrites event model data with the contents of the backup. Any changes that you have made to your event model subsequent to the time when the backup was saved are lost. Before you restore a backup, you may want to perform a backup of the current event model. See "Managing Backups" on page 256.

Note: The backup is immediately applied to the environment. Tealeaf recommends restoring backups during off-peak hours.

1. To restore a scheduled or on-demand backup, select **Tealeaf** > **Portal Management**.
2. In the Portal Management page, click the Backup & Restore link.
3. Click the Backup/Restore link.
4. For the backup you wish to restore, click the **Restore** link.
5. When the backup is uploaded, the contents of the backup are immediately applied to the environment.
6. Restart all Tealeaf services.

Deleting Backups

To delete a stored backup, click **Delete** next to the name on the Backup/Restore window.

Scenario - Migrating Events between Dev and Test

Optionally, this backup feature can be used to migrate complete sets of event definitions between Tealeaf environments. For example, if you have been working on event definitions in a Development environment, you can use this backup feature to migrate to your Test environment.

Note:

1. This backup and restore feature replaces entire sets of event and event objects. Objects in the destination system that do not exist in the source system are deleted. Before you begin, you should deploy any individual changes that you wish to keep from the destination system into the source system, if they do not exist there.
2. The license key in the two systems must be identical, or the event definitions will not be restored in the destination system. See "Managing Your Tealeaf License Key" on page 366.
3. In the following scenario, the event model from the Development system is backed up and restored into a Test system.
4. Note that this backup procedure overwrites event definitions. Depending on how event definitions are managed in a Production environment, this workflow may not be suitable for publishing into your live transaction stream.

1. Create an event model backup on the Dev system. See "Creating On-Demand Backups" on page 256.

2. Install a fresh Test system or validate that the existing Test system has the same license key as the Dev system.
3. Copy over the backup file that you created on the Dev system to the directory where backups are stored on the Test system. Copy the exported `.tlb` file to the import directory on the destination system.
4. Restore through the Portal on the Test system. See “Restoring Backups” on page 256.
5. The Test system now has the events and dimensions from the backup created on the Dev system, including dimension whitelist values and group list values.

Logging

You can check the status of scheduled backup jobs through the Data Service log. Scan the log files for messages that include `CreateBackup`.

- See "Configuring the Tealeaf Data Service" in the *IBM Tealeaf CX Configuration Manual*.

Backing Up Tealeaf

Tealeaf provides multiple mechanisms for backing up your Tealeaf system and your captured data.

Backup Type
Description

Event Model data
Through the Portal Management page, you can execute scheduled and on-demand backups of event model data. For more information, see Event Model Backup and Restore.

Database data
Using the Tealeaf Database Manager or SQL Server Management Studio, you can perform full backups of all Tealeaf databases.

- For more information on database backups in general, see "Database Backup Strategy" in the *IBM Tealeaf Databases Guide*.
- For more information on performing backups through the Tealeaf Database Manager, see "Using Tealeaf Database Manager" in the *IBM Tealeaf Databases Guide*.

Canister data
Tealeaf provides a pair of utilities for backing up session and index data stored on Tealeaf Canisters. For more information, see “TLBackup and TLRestore” on page 346.

scorecards/dashboards/reports
Through the Portal, you can import and export templates for scorecards, dashboards, and reports. For more information, see "Importing and Exporting Reports" in the *IBM Tealeaf cxView User Manual*.

Chapter 4. Tealeaf Management System

The Tealeaf Management System (TMS) centralizes configuration management for Tealeaf systems on Windows-based hosts.

Note: This version of TMS cannot be used to configure a IBM Tealeaf CX Passive Capture Application server, which is hosted in a non-Windows environment. PCA does provide a Web Console for managing configuration. See "Passive Capture Configuration via Web Console" in the *IBM Tealeaf Passive Capture Application Manual*.

In many customer environments, the Tealeaf solution is deployed across multiple servers, which may be spread across multiple data centers with varying access permissions and security requirements. As a result, applying configuration changes across the Tealeaf solution can be problematic, if you do not use TMS. Through a central configuration interface, the Tealeaf Management System enables the following Tealeaf configuration features:

- **Security:** TMS eliminates the need to have operating system-level access to Tealeaf machines for supported configuration changes.
- **Visibility:** Through TMS, you can review the services, settings, and pipelines available on each server where Tealeaf software is installed.
- **Complexity and Scaling:** Through the central interface, you can easily apply configurations across multiple servers, including ones that have been added after base installation of the solution.
- **Updates:** As needed, you can use the Advanced features of TMS to get and put files within the Tealeaf install directory on individual servers.
- **Data Segmentation:** Through TMS, you can configure access to session data through user groups.
  - Data Segmentation is managed through Search Server configuration within TMS. See "Configuring the Search Server" in the *IBM Tealeaf CX Configuration Manual*.

For information about accessing TMS, see *IBM Tealeaf CX Configuration Manual*.

For information about resolving any issues with TMS, see *IBM Tealeaf CX Troubleshooting Guide*.

Overview of TMS

TMS consists of a **master server** (normally the Portal server) and multiple **slave servers**, which are all other servers running Tealeaf software, with the exception of Passive Capture and cookie injectors. The TMS master server communicates configuration information to the appropriate slave servers. Each slave then updates the appropriate configuration file or registry.

Changes can be made from the Portal and scheduled for later updating on the relevant computers. Additionally, configuration settings can be defined as common to multiple servers, reducing the overall workload with multiple server systems.

- For more information on TMS configuration options, see the TMS Reference appendix.

© Copyright IBM Corp. 1999, 2014

Each of these TMS updates is called a **job**, which can be created, scheduled, rescheduled, and canceled through the Portal interface. Each job contains one or more **tasks**. A task implements the configuration changes, which allows for a set of related changes to be made at one time.

For more information on configuring the master TMS server, see “Configuring the TMS Server” on page 314.

External Configuration Management Tools

Before Release 7.0, Tealeaf configuration management was handled through a series of external tools, configuration files, and Windows registry settings. Many of these tools did not have useful features such as data entry validation and configuration data backup. Additionally, changes had to be made on the server hosting each component to modify.

The introduction of TMS in Release 7.0 began the process of centralizing configuration management through a single area in the Portal and providing a more robust system for making one change and deploying it across multiple servers.

Note: Over time, external configuration tools and resources may be deprecated. If a Tealeaf component can be configured and updated through TMS, you should get in the habit of using TMS to do so.

These external configuration programs are described elsewhere in the “cxImpact Administration Manual” on page xi and in the "Tealeaf CX Configuration Manual" in the *IBM Tealeaf CX Configuration Manual*.

Accessing TMS

To access TMS:

1. Login to the Portal as an admin user.
2. From the Portal menu, select **Tealeaf** > **TMS**.

- If you are experiencing difficulties with TMS, please see "Troubleshooting - Tealeaf Management System" in the *IBM Tealeaf Troubleshooting Guide*.

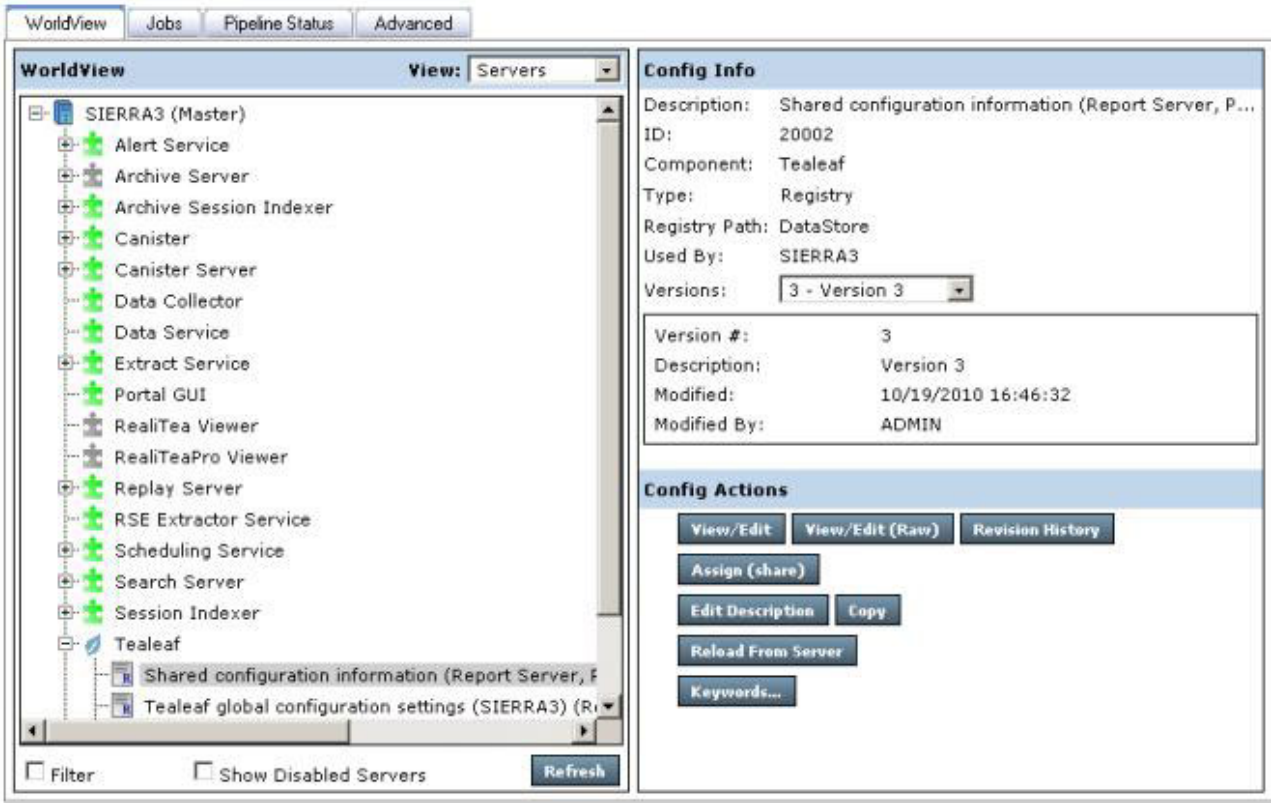

Figure 70. Tealeaf Management System

TMS Tabs

- **"TMS WorldView Tab" on page 262** - Displays all servers, components, and configurations in the TMS environment. You can view a different type of TMS item and navigate the tree to select the item to review, edit, assign, and more. See "TMS WorldView Tab" on page 262.
- **Chapter 5, "TMS Jobs Tab," on page 291** - Define and schedule TMS jobs, which consist of one or more tasks to start/stop/restart components and push configurations to slave servers. See Chapter 5, "TMS Jobs Tab," on page 291.
- **Chapter 6, "TMS Pipeline Status Tab," on page 297** - For each server in the environment, you can review the current status of its pipelines, including the data volumes and processing rates of each session agent. See Chapter 6, "TMS Pipeline Status Tab," on page 297.
- **Chapter 7, "TMS Advanced Tab," on page 303** - Backup and restore TMS configurations or export a summary for offline review. You can also scan all servers to discover new items that can be managed under TMS. See Chapter 7, "TMS Advanced Tab," on page 303.

Note: Tooltips are available for tabular configuration items.

Pipeline Editor

Through TMS, you can create, configure, and define the session agents in your Tealeaf pipelines through an easy-to-use graphical interface. See Chapter 8, "TMS Pipeline Editor," on page 307.

TMS Administration

TMS enables you to configure the TMS master server and other tasks through the interface. See “Configuring the TMS Server” on page 314.

TMS Reporting

TMS status messages over the past 24 hours can be included in the Tealeaf Status report. See “Tealeaf Status Report” on page 231.

See “Reporting for Tealeaf Administrators” on page 84.

TMS WorldView Tab

The TMS WorldView tab provides a top-down perspective on the servers, components, and configurations that the Tealeaf Management System manages in your Tealeaf environment. Through the WorldView tab, you can edit, copy, assign, and review the history of all configurations currently managed by TMS.

Selecting View

In TMS WorldView tab, you can display TMS items by Server, Component, or Configuration in the left pane.

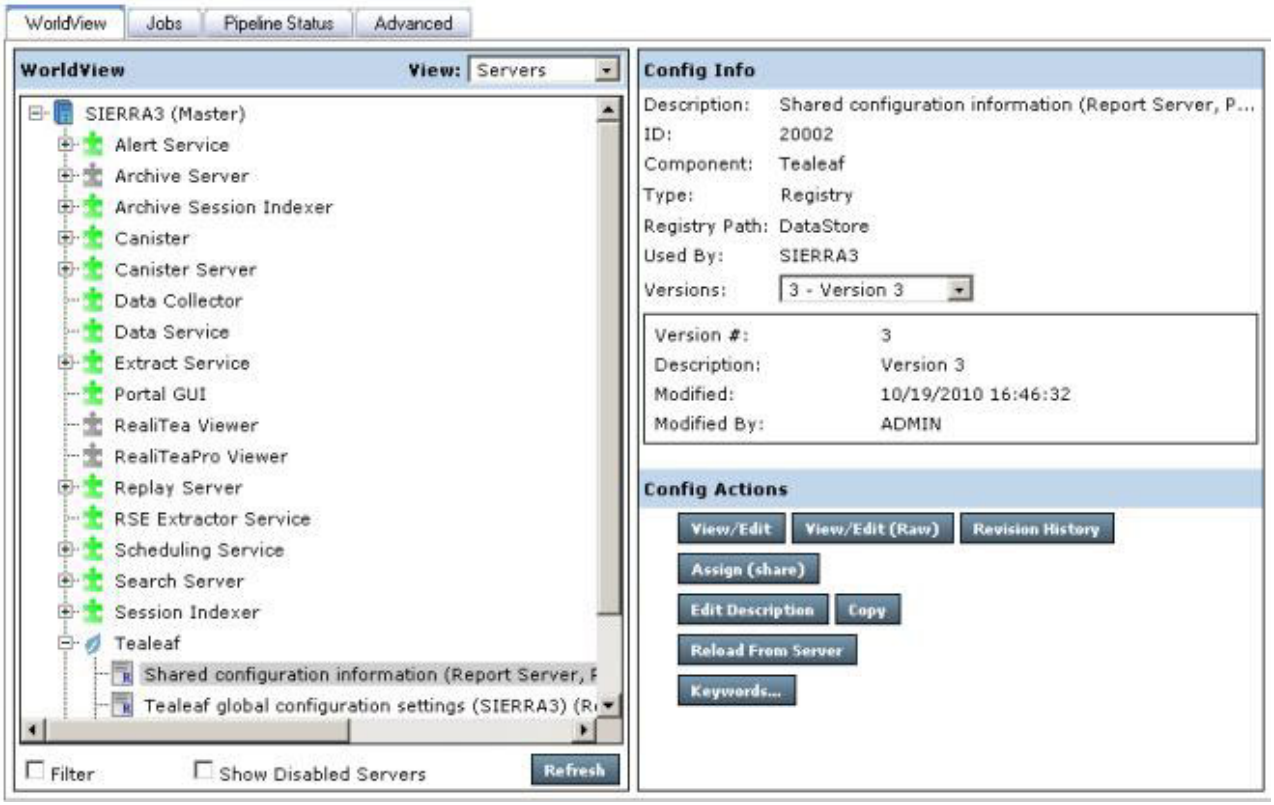

Figure 71. TMS WorldView Tab

To select a different view, use the View drop-down:

View **Description**

Servers
The individual computer systems. Click the sign to display the components installed on the selected system. See “Servers View” on page 266.

Components
The list of the individual Tealeaf subsystem components. Drilling down shows the servers with the component installed on them. See “Components View” on page 267.

Configs
Lists all Tealeaf components. Drilling down on any component shows you all of the configuration instances for that component. See “Configs View” on page 268.

To review items in the WorldView, expand a node to display the items beneath it. When you select an item, the available actions for the item are displayed in the Actions panel on the right side of the screen.

The WorldView display can be filtered using keywords assigned to servers, components and/or configurations.

- **Keyword filters:** - To add a keyword to an item, select the item and click **Keywords...** in the Actions panel. See “Managing Keywords” on page 265.
- **Display filters:** - To filter the WorldView display, select the Filter checkbox. Enter the desired keywords for the filter, and select the Enable Filter checkbox to turn on the filter. Click **OK**. The WorldView displays only the items containing the text filter. See “Applying Display Filters.”

Through TMS, you can disable individual servers. To display the servers that have been disabled, select the Show Disabled Servers checkbox. When it is selected, all enabled and disabled servers are displayed.

- To enable or disable a server, select the Servers view. Then, select the server. To disable an active server, click **Disable Server**.

  Note: You can disable any server except the TMS master server.
- If disabled servers are displayed, then when you select one, you can enable it by clicking **Enable Server**.
- To refresh the WorldView display, click **Refresh**.

Applying Display Filters

Through the Edit Filter dialog, you can define filters to view a subset of items in the WorldView.

- To edit filters, click the Filter checkbox at the bottom of the WorldView tab.

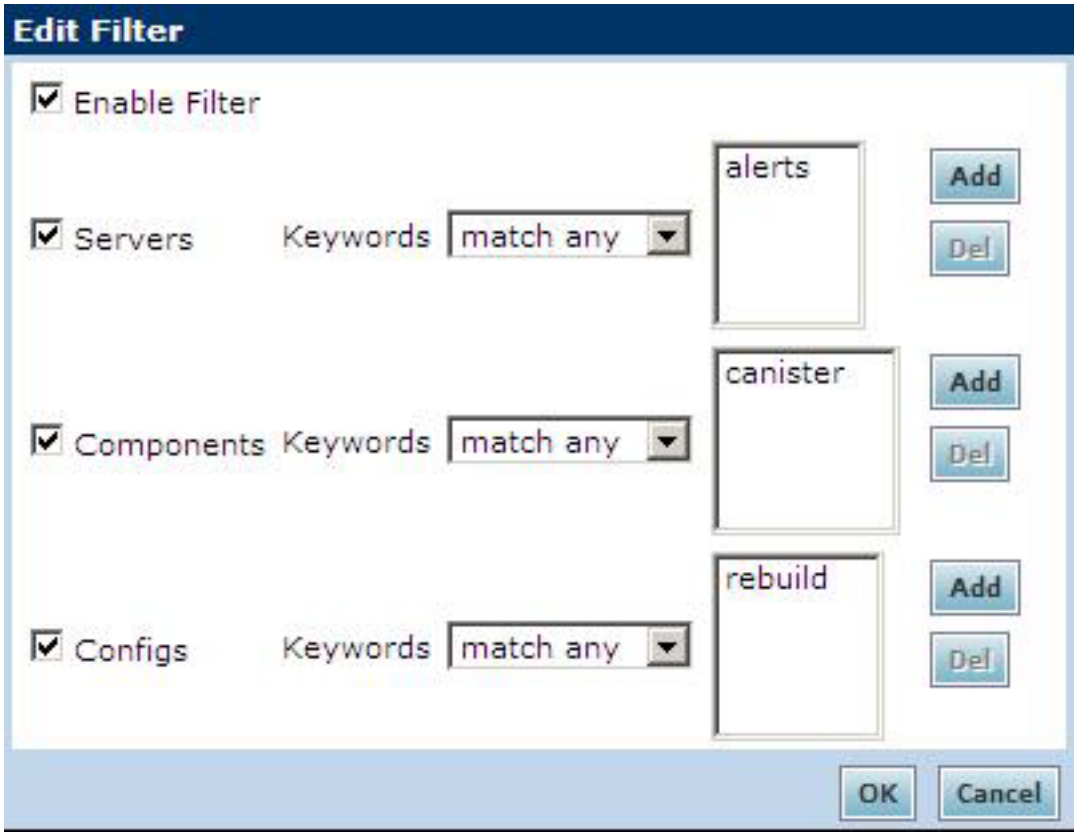

Figure 72. Edit Filter

To add a filter:

1. To apply filtering, click the Enable Filter checkbox. The filters specified in the dialog are applied to all views in the WorldView.
2. Select the one or more types of items to filter (Servers, Components and/or Configs).
3. Select the matching condition from the Keyword drop-down:
   - `match any` - Display any TMS item that matches any of the selected keywords.
   - `match all` - Display any TMS item that matches all of the selected keywords.
   - `match none` - Display any TMS item that does not match any of the selected keywords.
4. To specify matching keywords, click **Add**. Select the keyword from the list of available ones in the drop-down. Then, click **OK**.
   - To delete a keyword from a view filter, click the keyword, and then click **Del**.
   - To change keywords, select **Keywords...** in the Actions panel. See "Managing Keywords" on page 265.
5. To apply the selected filters, click **OK**.

After keywords are applied and the filter is applied, only items of the selected type that meet the filter criteria are displayed. All other items are hidden.

- Applying the filter does not change your current view.

- Applying the filter only filters on the view types you specify. If you filter on Configs and then display Components view, all items are displayed.

Managing Keywords

Through TMS, you can group servers, components, and configurations by assigning keywords. When a group of TMS items is assigned to a keyword, you can filter the display so that only those items are displayed, which enables rapid review and updating of related configuration items.

Note: Keywords apply to individual instances of components. When the WorldView tab is in Config view, the display reflects types of components. So, keywords do not apply.

To review, add, or remove keywords, select a TMS item, and then click **Keywords...** in the Action panel. The Keywords dialog is displayed. From the Keywords dialog, you can also click **Remove Keyword from System**. This removes the keyword not only from the server, component, or config selected, but also from the entire system.

In the Keywords dialog, all keywords that are assigned to the current item are displayed in the left panel.

- To add a keyword, click **Add Keyword**. See "Adding keywords."
- To remove a keyword from the current item, select the keyword and click **Remove Keyword**.
- To save changes, click **OK**.

Adding keywords

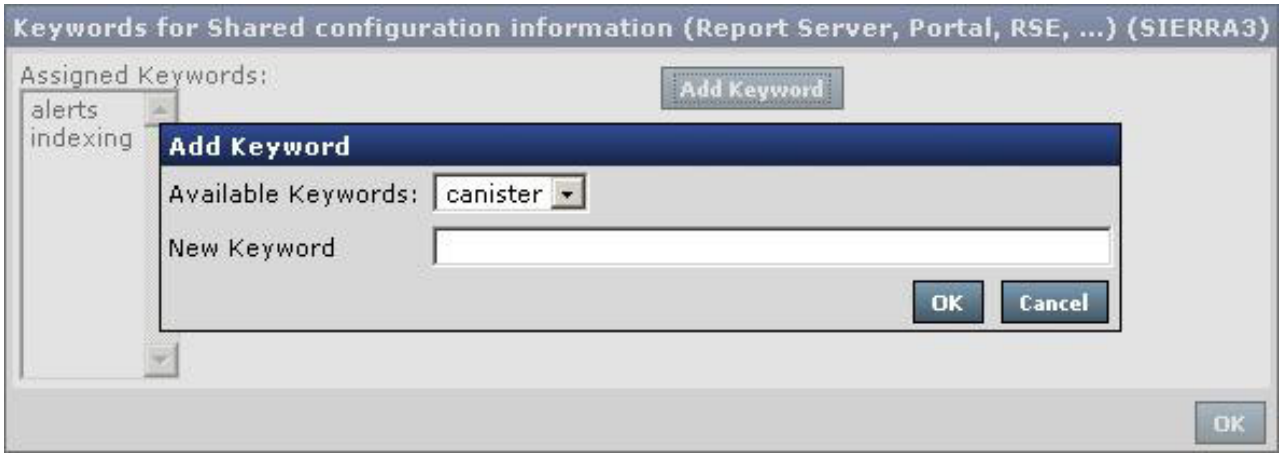

Figure 73. Adding keywords

In the Adding keywords dialog, you can enter a new keyword to use or select from the list of available ones. This list contains all keywords that have not been assigned to the current item.

- Adding a new keyword automatically assigns the keyword to the current item.
- To save changes, click **OK**.

Importing and exporting templates

From the Worldview Tab in the Tealeaf Management System, you can import and export templates to the Replay Server to customize your native app session replay.

Templates are installed under `ReplayServer\TLT` directory relative to IBM Tealeaf install. However, access to the Replay Server might be restricted. If so, you can you **Import Templates** and **Export Templates** in TMS to make updates to your native app session replay templates.

Click **Import Templates** to choose from a local TAR file, which bundles the templates. The structure of the TAR file should only contain a TLT folder that consists of all templates. Subfolders should be in the TLT folder. If the TAR file does not comply, it is not accepted through the import function. Once the TAR file is validated, the existing TLT folder is moved to the Backup folder under Replay Server and renamed with the current timestamp. The new folder is namd with the TLT name and the current timestamp.

Note: Only TAR files that are created using the 7-zip utility are supported.

Click **Export Templates** to create a TAR from the current TLT folder and download it on your local system.

For more information about customizing native app session replay templates, see "Native app session replay customization" in the *IBM TealeafCX Configuration Manual*.

Servers View

This view lists all Tealeaf servers in the TMS environment.

- A blue icon next to a server name indicates the TMS master server, while a gray icon indicates a TMS slave. See "TMS Master Server."
- A Tealeaf icon indicates the Tealeaf component, which holds common configuration items. It is present and running on all TMS master servers. See "Tealeaf Component" on page 277.
- To display a list of components installed on a server, click the icon next to the server name. See "Components View" on page 267.

Expanding a component node displays a list of configurations for that component assigned to the selected server.

- See "Configs View" on page 268.
- If you select the Transport Service or Event Bus configurations, you edit them in the Pipeline Editor. See Chapter 8, "TMS Pipeline Editor," on page 307.

Server Actions

See "Config Actions" on page 268.

TMS Master Server

The TMS master server is indicated with a blue icon next to it, followed by the server name and `(Master)`.

Note: You cannot enable or disable the TMS master server.

- To view basic information about the server, click its name.

In the Server Info panel, you can review information about the server, including the Tealeaf roles assigned to it. Tealeaf roles correspond to Tealeaf services or servers installed on the indicated server.

Components View

The Components view displays a list of the Tealeaf components installed on one or more of the servers. To display a list of component instances, click the icon next to the component instance.

- A green icon next to a component instance indicates that it is currently running on that server.
- Component instances with a grey icon are in a stopped state.

To review basic information about the component instance, click it and then click **View/Edit**. See "Viewing and editing configurations" on page 269.

- If you select the Transport Service or Event Bus configurations, you edit them in the Pipeline Editor. See Chapter 8, "TMS Pipeline Editor," on page 307.

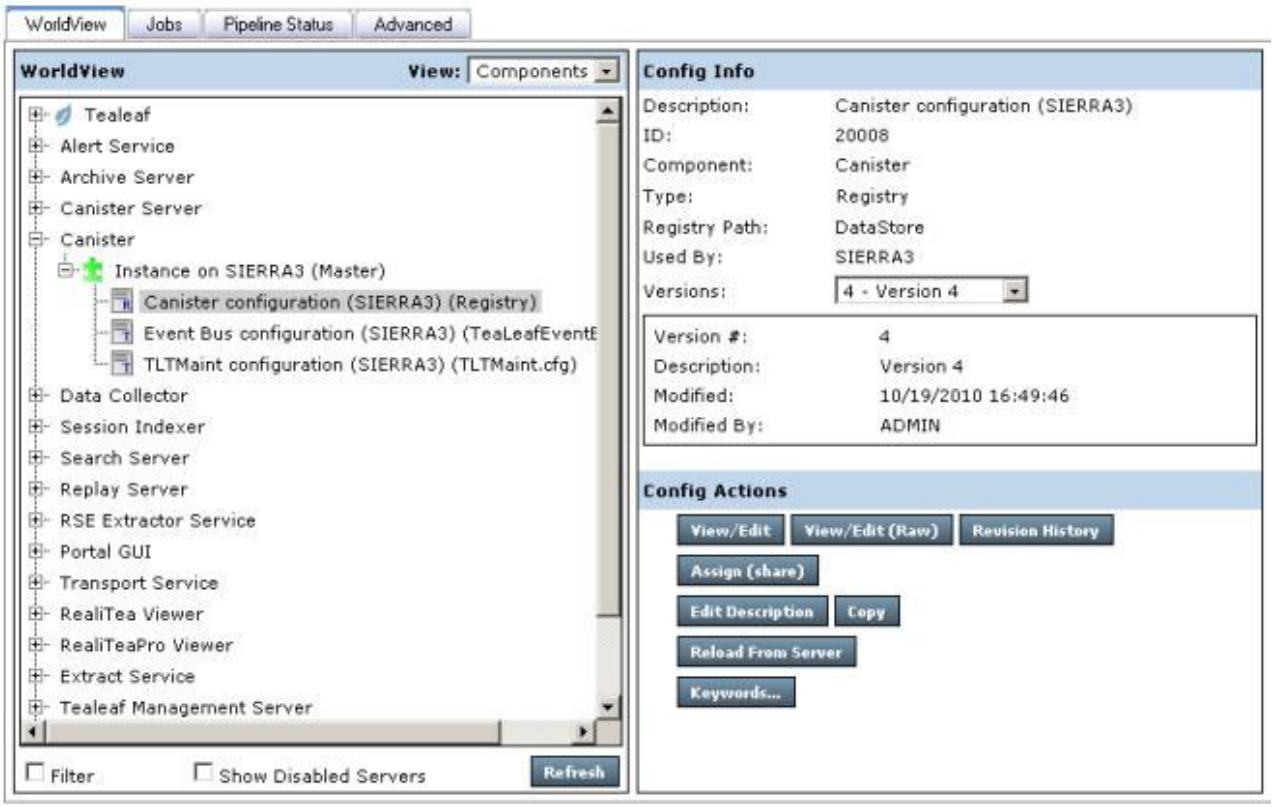

Figure 74. TMS World view

Component Actions

The following actions apply to components:

- **Start** - Start the component.
- **Stop** - Stop the component.
- **Restart** - Restart the component.

See "Component update behaviors" on page 276.

See “Config Actions.”

Configs View

Similar to the Components view, this view initially presents a list of components. When you expand a component node, a list of all configurations associated with that component for all servers is displayed.

Note: This view is the only place where you can see configurations that are not currently assigned to a server.

- Any component without a + next to the icon either doesn't have any configurations associated with it, or there are none currently assigned for that server.
- A yellow icon next to a configuration indicates that the configuration has been modified but hasn't yet been pushed to the server; the server's actual configuration file or registry hasn't been updated with the new values.
- Configuration icons with a `T` in them are managed through a configuration file. The configuration filename is listed after config name and server.
- Configuration icons with an `R` in them are managed through the Windows registry.
- Configuration icons with a red + in them are shared by more than one server.
- If you select the Transport Service or Event Bus configurations, you edit them in the Pipeline Editor. See Chapter 8, “TMS Pipeline Editor,” on page 307.

To view configuration information, click the configuration name and then click **View/Edit**. See “Viewing and editing configurations” on page 269.

Config Actions

Action buttons appear only when they are applicable to the currently selected item.

- **View/Edit** - Review and edit the configuration. See “Viewing and editing configurations” on page 269.
- **View/Edit (Raw)** - Review and edit a text-based configuration in its raw format. See “Editing raw configurations” on page 271.
- **Revision History** - Review the revision history for the selected item. See “Revision history” on page 273.
- **Assign (share)** - Assign the configuration to other TMS servers. See “Assigning configurations” on page 272.
- **Edit Description** - Edit the description for the configuration, which appears in the Config Info panel.
- **Change Filename** - For file-based configurations, you can change the name of the file. See “Editing configuration filenames” on page 275.
- **Copy** - Copy the configuration. See “Copying configurations” on page 275.
- **Reload from Server** - Reload the configuration from the selected server.
- For some configurations, you can choose to encrypt them on disk in the Config Info panel. See “Encrypting configurations on disk” on page 275.
- **Keywords** - Add or edit keywords associated with the configuration. See “Managing Keywords” on page 265.
- **Rename** - For file based configurations, you can rename the file stored in the server. See “Renaming configurations” on page 273.

Viewing and editing configurations

You can view or edit any saved version of a configuration.

To select a version of the configuration to view or edit:

1. Select the configuration.
2. In the Config Info panel, select the version from the drop-down.
3. Click **View/Edit** in the Config Actions panel.
   - To view the configuration in raw format, click **View/Edit (Raw)**. See "Editing raw configurations" on page 271.
4. If TMS detects differences between the version stored in the TMS datastore and the version stored on the server, you can choose to overwrite the TMS version with the server-stored version. See "Externally modified configurations" on page 270.

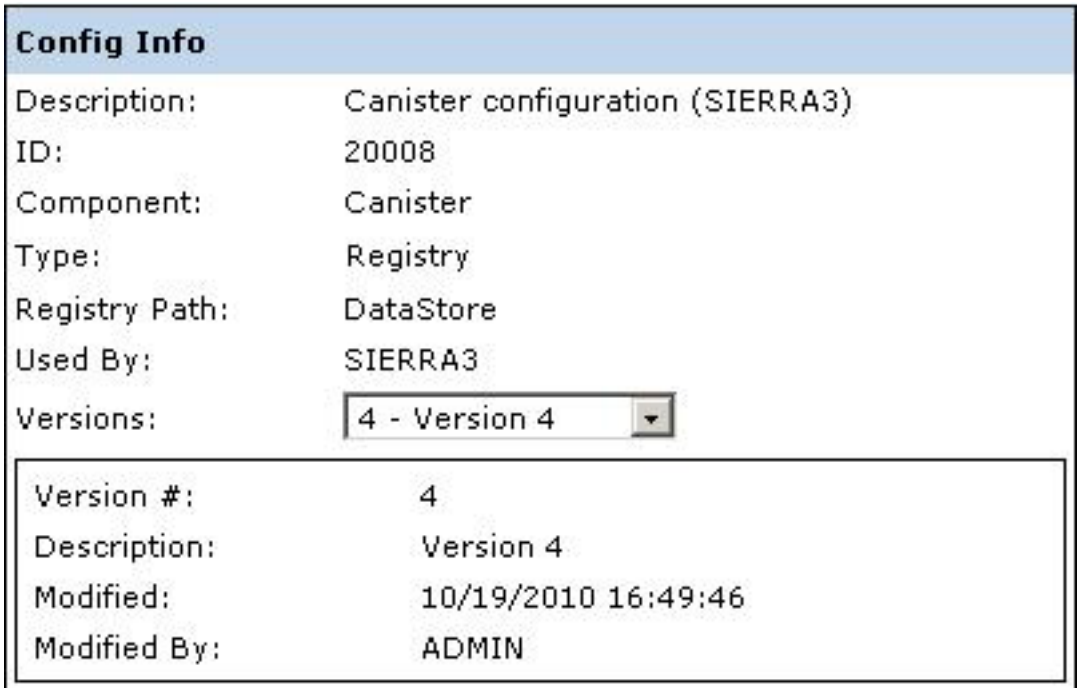

Figure 75. View configuration version

- TMS retains each edited version of a configuration. To view a different version of the configuration, select the version from the Versions drop-down.
- To load a configuration that has been modified outside of TMS (e.g. using one of the external configuration programs), click **Reload from Server**. This updates the TMS datastore with the configuration data found on the specified server (if modified).

Option Description

`Source Server`
Select the server from which to load the configuration.

`Version Label`
If desired, you can enter your own version label in the space provided. This label appears in the Version drop-down.

Share Local Data With All Listed Servers
When this option is selected, TMS pushes the local version that you upload to all configured consumers of this configuration.

Externally modified configurations

If you have added or modified configuration settings in configuration files local to the currently selected server, you may be prompted to choose to load an externally modified version:

```
Configuration was modified outside of TMS on the following server(s):
MYSERVER
Reload configuration from externally modified version?
```

If you choose **Yes**, the local configuration file becomes the latest version of the configuration file stored for the current server.

- Changes to local registry settings also result in this prompt if the settings appear in the standard configuration in TMS.

  Note: Making changes to the local registry settings is not recommended.
- If the configuration loaded from the local server is to be shared among multiple servers, you must manually create a task to publish the changes to the other servers. See Chapter 5, "TMS Jobs Tab," on page 291.

Editing configurations

In the configuration edit dialog, you can see a Version Description field. If you make any changes to the configuration, you can enter a description for the new version here.

The actual configuration data is displayed in a text edit window, as a list of config items that can be individually edited, as in a configuration-specific interface such as the Pipeline Editor. The method of editing is determined by the type of configuration.

- For configurations that display in a text edit window (the words `Config Text:` appear above the edit window), you can edit the configuration data as you normally would with a text editor.
- For configurations that are displayed in the Configuration Editor (the words `Config Sections:` appear above the configuration edit window), you can review and edit the config values in one or more config sections.
  - The section contents can be displayed or hidden by clicking on the section name (or the +/- next to it).
  - To edit a config item value, click the item. A small dialog includes a brief description of the item. Data entry is limited to the correct data type (number, text, boolean, etc.) and range.

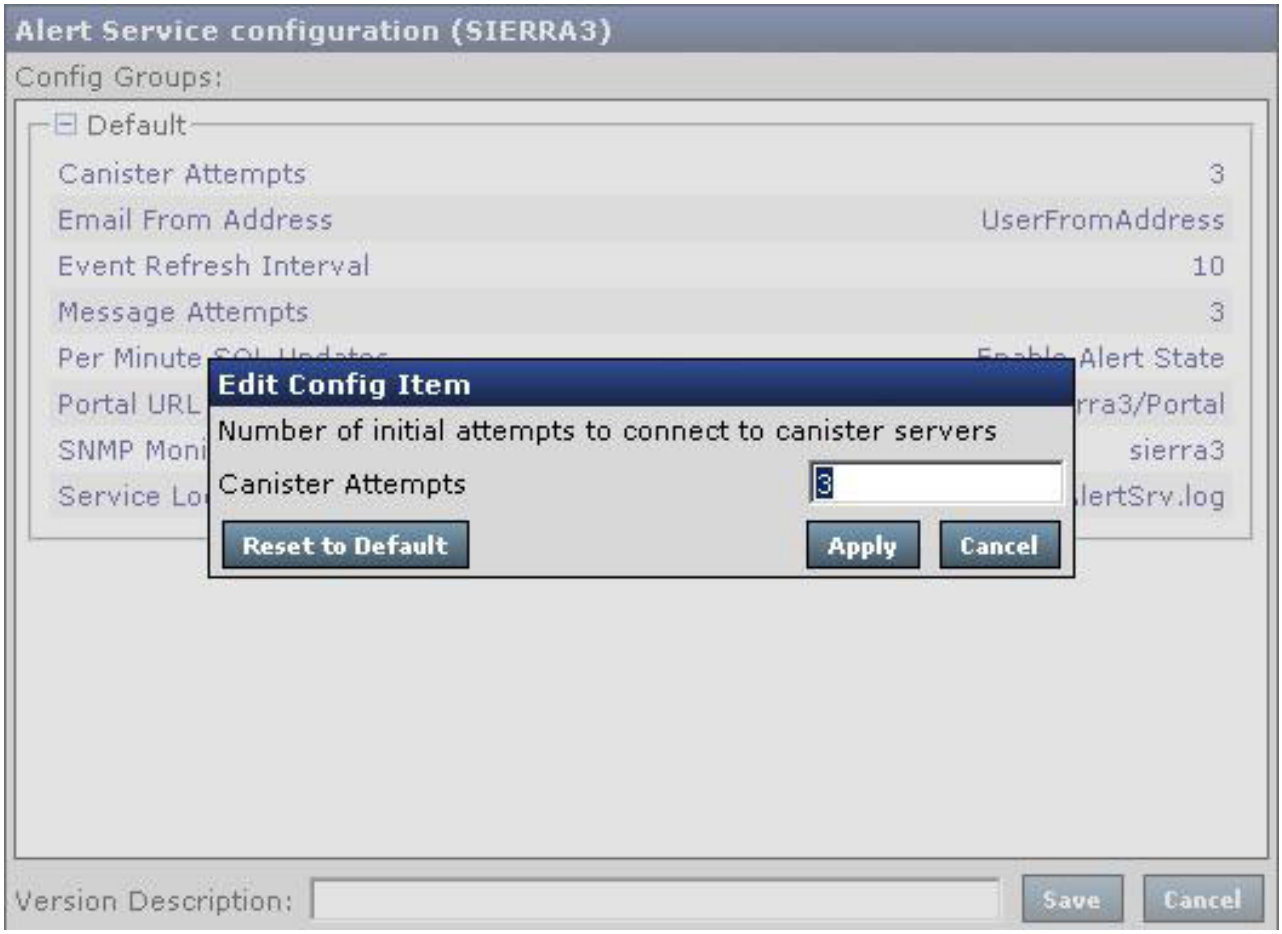

Figure 76. Configuration window

To save changes made to a configuration, enter a version description and click **Save**. The configuration is saved as a new version in the master TMS datastore.

After saving, you can choose to add tasks to push the new version to all servers that use it. Click **OK** to create a new TMS job composed of one or more tasks to push the configurations to the appropriate destinations.

- For tasks that push configurations, you can choose to force the affected components to restart after the new configuration has been applied.
  - The Push/Restart job is automatically submitted when you click **Add Tasks and Submit**.
  - To submit the job manually at a later time, you can click **Add to Current Job**.
  - To submit the configured job to TMS, make any changes to the job and then click **Submit**. See Chapter 5, "TMS Jobs Tab," on page 291 below.

Editing raw configurations

In TMS, the most useful configuration items are exposed in the standard format. Optionally, you can review the raw configuration, which contains all settings associated with the configured item.

Note: Avoid editing raw configuration data, as no validation of entries is performed.

- To edit a raw configuration, select the item in the WorldView tab and click **View/Edit (Raw)** in the Config Actions panel.

The raw configuration exposes all settings contained in the TMS data dictionary for the selected item. In addition to the items available in the standard configuration, the raw configuration includes:

- read-only items that cannot be changed by users
- hidden items retained for debugging purposes

Note: The items in the standard version of each configuration are documented and exposed for your use. Unless directed by Tealeaf to change them, it is recommended that you avoid making changes to hidden configuration items.

Items that are hidden and read-only may include obsolete settings, which are maintained for upgrade purposes.

Assigning configurations

A configuration is not actively in use unless it has been assigned to a server. When a server is added to TMS, all current configurations are added and assigned to that server.

1. To assign a configuration to one or more servers, select the desired configuration and then click **Assign**. A list of servers to which the configuration can be assigned is displayed.
2. Click the checkbox next to each server name to assign or un-assign the configuration, and then click **Assign** to complete the assignment.
   - To cancel the assignment, click **Cancel**.
3. The configuration is assigned immediately to the selected servers. Those servers are updated the next time the corresponding Push Configuration job is executed by TMS.
   - For more information on scheduling TMS jobs, see Chapter 5, "TMS Jobs Tab," on page 291.

You can share a single configuration, assuming there are no server-specific config values among multiple servers.

Note: A server and component cannot have multiple assigned configurations with the same name. Assigning a configuration to a server that already has a configuration of that type assigned causes the original configuration to be unassigned from that server.

Un-assigning a configuration

If you have accidentally assigned a configuration to a server and wish to revert to another version, please complete the following steps.

1. In the WorldView tab, select `Configs` from the View drop-down.
2. Select the component that contains the configuration you wish to use.
3. Navigate the tree to select the configuration to use.
4. In the Config Actions pane, click **Assign (share)**.
5. Select the server or servers to which to assign the configuration, replacing the one you originally assigned.
6. Click **Assign**.
7. Submit a task to complete the change.

Renaming configurations

Through TMS, you can rename the configuration as it appears in the Portal and in the Windows registry. To rename a configuration, select it and click **Rename**. In the Description field, enter a new name for the configuration, and click **Save**. The configuration name is changed.

Revision history

In the Revision History window, you can review the history of changes to the selected configuration.

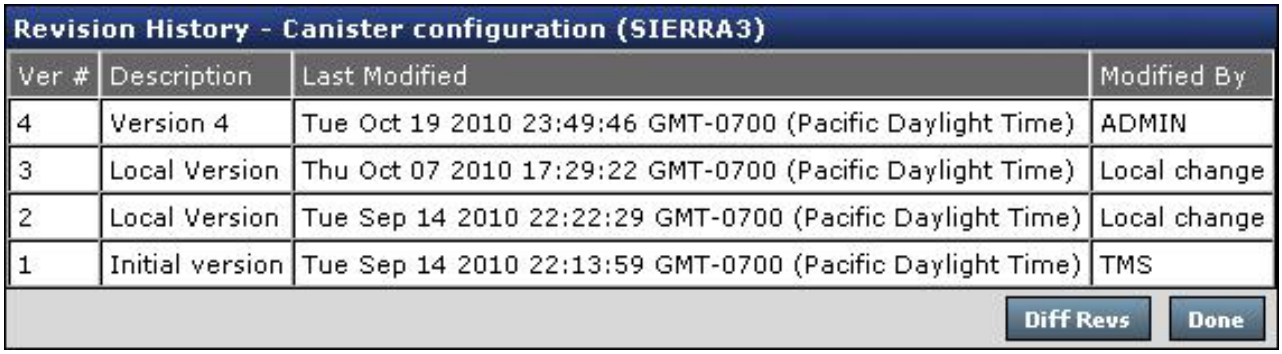

Figure 77. TMS configuration revision history

- To compare differences between two changes, select the two revisions. Then, click **Diff Revs**. See "Revision differences."
- To close, click **Done**.

Column
Description

Ver # Version number

Description
The description of the change

Last Modified
The date and timestamp for when the modification was committed

Modified By
The Tealeaf user account through which the change was committed.

- TMS - changes modified by the TMS account are system changes.

Revision differences

In the Revision Differences window, you can review the differences between the two selected revisions.

- To view differences in a single view, click the Unified View tab.
- To see the changes side-by-side in separate panes, click the Side-By-Side View tab.

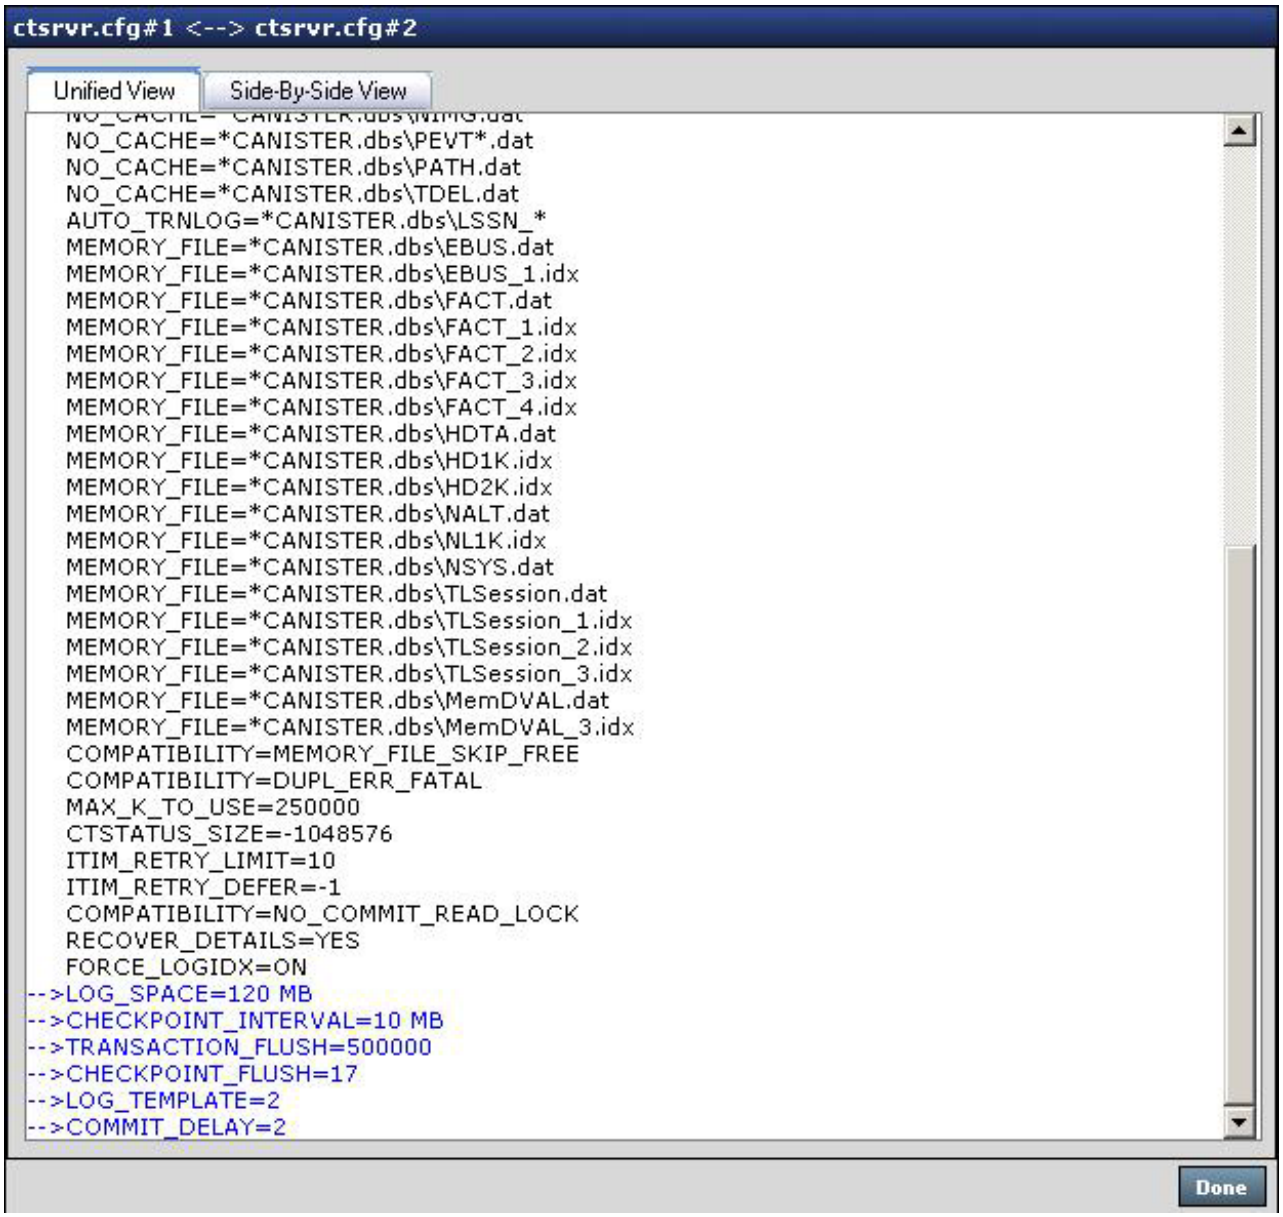

Figure 78. TMS configuration revision differences

Unified view differences

Change
Identifier

Added lines
blue text

Deleted lines
red text

Updated lines
old line: red text new line: blue text

Side-by-side view differences

Change
Identifier

Added lines
earlier version: gray background later version: blue background

Deleted lines
earlier version: red background later version: gray background

Updated lines
earlier version: red background later version: blue blackground

Editing configuration filenames

For file-based configurations, you can rename the file in which it is stored.

Note: You should rename a configuration file only if you understand the impacts of the name change. Some configuration files require a specific name. For example, `TeaLeafCaptureSocket.cfg` is the required filename for the Transport Service configuration.

1. To rename a file-based configuration, select it and click **Change Filename**.
2. In the Filename field, enter a new filename for the configuration.
3. Verify that you have included an appropriate extension.
4. Click **Save**.
5. The configuration filename is changed on the server and in TMS.

Copying configurations

You can copy a configuration to use it as a basis for experimentation or developing related configurations.

To copy a configuration:

1. Select the configuration.
2. Click **Copy**.
3. In the Description field, enter the name of the new configuration as it appears in TMS.
4. For the initial version of the copied configuration, you can enter an appropriate description to indicate that it was copied from another description.
5. To save the copy, click **Save**.
6. The configuration appears in TMS.

This new configuration can be assigned independently of the original configuration.

Encrypting configurations on disk

For some configurations, such as Transport Service and Privacy Rules, you can choose to encrypt the configuration file on the disk where it is stored.

Note: By default, Tealeaf does not encrypt configurations. There is no impact on performance to use encryption.

To encrypt a configuration:

1. Select the configuration to encrypt.

2. If the configuration can be encrypted, click **Encrypt on Disk**, which appears in Config Info panel.
3. The Update Servers dialog appears. Make your selections. See "Update Servers."
   - To remove encryption from a configuration, select it and click **Remove Encryption** in the Config Info panel.

Update Servers

In the Update Servers dialog, you can specify whether to update servers affected by the current change.

Figure 79. Update Servers dialog

- **Add Push Config Task(s)** - Select this item to push the configuration changes as one of the tasks of the job.
- **Add Restart Task(s)** - Select this item to add a server restart as a task. If a Push Config task has been added to the job, this item is executed after it. See "Component update behaviors."
- **Servers** - Select the servers in the TMS environment to receive this job.
- To add the selected tasks to the job currently selected in the Jobs tab, click **Add to Current Job**.
- To create a new job with the selected tasks and to execute it immediately, click **Add Tasks and Submit**.
- To cancel server update, click **Cancel**. The changes are still saved to the local configuration(s).

Component update behaviors

When TMS updates a server or service, TMS checks the service for dependencies on other services and for services that are dependent on the restarted one.

Based on those dependencies, the following restart behaviors occur:

- If the service has dependents, only the specified service is restarted. Dependents are not restarted.
- Stopping a service with dependents stops service and all dependent service(s).
- Restarting service with dependents restarts service and any dependents that are currently running.

Note: Restarting some components forces a restart of the IIS web server, which disconnects the Portal application from the Tealeaf Management Server and generates a `service unavailable` error in this page. When IIS completes its restart, refresh this page through the browser, and subsequent TMS operations should behave normally.

Refreshing the display

To refresh the WorldView, click **Refresh** at the bottom of the panel.

Tealeaf Component

Under each TMS master and slave server, the Tealeaf Component contains global settings and configurations that are shared by multiple components.

The Tealeaf component is a container component and never shows as running.

From the Tealeaf component, you can **Start**, **Stop**, or **Restart** services for all components. When you click **Start**, **Stop**, or **Restart**, you see a confirmation dialog. Once you confirm the action, you are logged out of the Portal. You can also cancel the action.

You can add and remove keywords from the Tealeaf Component.

Note:

1. If you are making changes to a configuration in a Tealeaf Component that applies across all TMS servers, then you must apply the change on each TMS server.
2. After restarting the Tealeaf component on the server hosting the Portal application, you must refresh your browser page to see the changes.

Tealeaf global configuration settings

Setting Description

`Days to Keep Log Files`
Number of days to retain log files on the server. For more information, see the Data Retention section of “CX Settings” on page 192.

`Roll Time Zone`
The time zone recognized across all Tealeaf servers when the LSSN files for each canister are rolled to begin storing a new day of data in a new LSSN file.

Note: This value must be set to the same time zone for all Tealeaf canisters in the environment. In multi-server environments, failure to assign each server to the same `Roll Time Zone` value can generate incorrect results. For more information, see "Configuring the System Timezone" in the *IBM Tealeaf CX Configuration Manual*.

`TMS Port`
The port to use to connect to the local TMS server. The default value is `20000`.

`TeaLeaf Log Directory`
The location of the Tealeaf log directory on the server. The default value is `Logs` inside the Tealeaf install directory.

TeaLeaf Install Folder
The install folder for Tealeaf components on the server.

Tealeaf mail configuration settings

These settings describe how Tealeaf components on this server can email alerts.

Setting Description

Mail Server Name
The mail server to access.

Password
The password to use when accessing the server to send emails.

SMTP Port
The port to use on the mail server.

Use Anonymous Email
When set to `True`, the User Name is omitted from the sent mail.

User Name
The user name to use to access the email server.

- The above settings do not apply to the Tealeaf Portal, whose email capabilities are configured in a different area. See "CX Settings" on page 192.

Shared configuration information

These settings enable configuration of TMS reporting capabilities for the Portal, IBM Tealeaf cxResults, and the Tealeaf system database.

Portal settings

These settings apply to the Portal application associated with this server.

Setting Description

Portal Session Timeout
The number of minutes that a Portal session is allowed to remain inactive, after which the user is automatically logged out.

Date Format
Defines the format for dates entered or displayed in the Portal.

Authentication Method
This setting defines the encryption method used to manage authentication. Supported methods are the following: `3DES`, `RC2`, `NT`, or `MD5`.

Search Server Timeout
The number of minutes that a request from the Portal to Search Server is allowed to remain inactive, after which the request times out.

Allow Mixed Mode Login
When set to `True`, Tealeaf users can choose to login using Portal authentication or NT authentication in the Portal login screen.

Database settings

Note: For most changes to these database settings, a corresponding change is required within the SQL Server Administration Console.

Setting Description

Report Database Server
The host name for the Reports server.

- For a local database, enter the value `localhost`.
- If it is a name instance, enter the value in the following format: `<server_name>\<instance_ name>`.

System Database
The name of the Tealeaf System database. By default, this value is `TL_SYSTEM`.

cxImpact Reports Database
The name of the Report database. By default, this value is `TL_REPORTS`.

Session Segment Database
The name of the RSE database. By default, this value is `TL_RSEXTRACTOR`.

System Statistics Database
The name of the Tealeaf Statistics database. By default, this value is `TL_STATISTICS`.

Admin User Name
The user name for the admin user of the database. By default, this value is `TLADMIN`.

- This password for this account is automatically generated during the install process. To change the password for this account, click **Admin Database Password**. Enter the new value twice and click **OK**.

Portal User Name
The user name for the Portal user of the database. By default, this value is `TLUSER`.

- This password for this account is automatically generated during the install process. To change the password for this account, click **Admin Database Password**. Enter the new value twice and click **OK**.

Reports Database Server Port
The port to use to communicate with the Reports server. By default, this value is `1433`.

Log Level
Database logging level (1-9).

Note: This value is also used by the Portal to set logging levels.

Processing Interval
The interval in seconds between data collection runs. By default, this value is set to `300` (five minutes). When this value is set to `0`, the Reporting Service is executed once and then stops.

Connection Timeout
The maximum number of seconds permitted for the SQL Server database to respond before the connection is timed out. By default, this value is `600` seconds (five minutes).

IBM TealeafcxResults DB settings

These settings apply to Visitor database, which is part of IBM Tealeaf cxResults.

Note: This tab appears only if you have licensed and installed IBM Tealeaf cxResults. IBM Tealeaf cxResults is no longer available as a newly licensed product as of Release 8.7. Customers that licensed IBM Tealeaf cxResults in Release 8.6 and

earlier may continue to use and receive support for the product in Release 8.7 and later. For more information, please contact Tealeaf Customer Support .

Setting Description

Visitor Database Server
The server hosting the Visitor databases.

Visitor Database Port
The port to use to communicate with the Visitor databases.

Path for Visitor Database files
Path on the database host to the location of the files for the Visitor database.

Visitor Report Database
The name of the Visitor reporting database. By default, this value is `TL_VISREPORT`.

Visitor Staging Database
The name of the Visitor staging database. By default, this value is `TL_VISSTAGE`.

Component Configuration References

Use the links below to review documentation for the available configuration item. The following set of links is available in `Servers` view in the TMS WorldView tab.

- Alert Service
  - Alert Service configuration. See "Configuring the Alert Service" in the *IBM Tealeaf CX Configuration Manual*.
- Canister
  - Canister configuration. See "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.
  - Event Bus configuration. See "Tealeaf Event Bus" in the *IBM Tealeaf cxConnect for Data Analysis Administration Manual*.
  - TLTMaint configuration. TLTMaint rarely requires post-install configuration.
- Canister Server
  - Canister Server configuration. See "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.
- Data Collector. See “Data Aggregation and Retention” on page 241.
- Data Service. See "Configuring the Tealeaf Data Service" in the *IBM Tealeaf CX Configuration Manual*.
- Extract Service
  - Tealeaf Extract Service configuration. See "Configuring the Extract Service" in the *IBM Tealeaf CX Configuration Manual*.
- Portal GUI. No configuration required.
- IBM Tealeaf CX RealiTea Viewer. See "RealiTea Viewer Menu Tools" in the *IBM Tealeaf RealiTea Viewer User Manual*.
- IBM Tealeaf CX RealiTea Viewer Pro. See "RealiTea Viewer Menu Tools" in the *IBM Tealeaf RealiTea Viewer User Manual*.
- Replay Server
  - Replay Server configuration. See "Configuring the Replay Server" in the *IBM Tealeaf CX Configuration Manual*.
- RSE Extractor Service. See “RS Extractor Settings” on page 223.

- Scheduling Service
  - Scheduler Service configuration. See "Configuring the Scheduling Service" in the *IBM Tealeaf CX Configuration Manual*.
- Search Server
  - Search Server configuration. See "Configuring the Search Server" in the *IBM Tealeaf CX Configuration Manual*.
- Session Indexer
  - Index Service configuration. See "Configuring CX Indexing" in the *IBM Tealeaf CX Configuration Manual*.
- Tealeaf. See “Tealeaf Component” on page 277.
- Tealeaf Management Server
  - Tealeaf Management Server configuration. See “Configuring the TMS Server” on page 314.
- Transport Service. See "Configuring the Transport Service" in the *IBM Tealeaf CX Configuration Manual*.
  - DataParser Search Templates. See "Data Parser Session Agent" in the *IBM Tealeaf CX Configuration Manual*.
  - Privacy Filter configuration. See "Privacy Session Agent" in the *IBM Tealeaf CX Configuration Manual*.
  - Transport Service configuration. See Chapter 8, “TMS Pipeline Editor,” on page 307.

Proactive alerts

Proactive alerts provide health-based alerts for the following IBM Tealeaf components: Canister server, data services, and CX Passive Capture Application (CX PCA) server.

After configuring proactive alerts for a component, you receive an alert through an email. The email contains the following details:

- Parameter name for which the alert is generated
- Current value of the parameter
- Threshold value that is configured for the parameter
- Component name for which the parameter is configured

Keep the following points in mind when you configure proactive alerts:

- If an IBM Tealeaf component is scheduled to be shut down for maintenance, disable proactive alerts for the component.
- If proactive alerts feature for a component is disabled for its maintenance, manually enable proactive alerts for the component after the component is up and running, and restart the TMS server.
- If Canister servers or CX PCA servers are added or removed, or if any threshold values have changed, manually change the configuration of proactive alerts for the component, and restart the TMS server.

Important: Your changes are effective after you restart the TMS server.

Configuring proactive alerts for IBM Tealeaf components

You can use the Tealeaf Management System (TMS) to configure proactive alerts for Canister servers, data services, and PCA servers.

You can enable or disable proactive alerts by using the **Turn on Proactive Alerts** check box in the Proactive Alerts Configuration window. To disable proactive alerts, clear the **Turn on Proactive Alerts** check box. Once disabled, proactive alerts are not generated for any component even if you have configured an alert condition for the component.

Complete the following steps to configure Canister servers, data services, and CX PCA servers to receive proactive alerts:

1. On the TMS home page, go to the **Worldview** tab.
2. Expand your master server and then expand the **Tealeaf Management Server** node.
3. Click the **Proactive Alerts Configuration** *MasterServerName* link. The Proactive Alerts Configuration window opens.
4. Ensure that the **Turn on Proactive Alerts** check box is selected. If you want to turn off the Proactive Alerts feature, clear the **Turn on Proactive Alerts** check box.

   Important: If you have configured proactive alerts for any component, and if you clear the **Turn on Proactive Alerts** check box, proactive alerts are not generated for any component.
5. In the **Global PA Rules** area, click **Edit**. The PA Global Edit window opens. Refer to the following table to add information for email alerts:

Table 27. Information about the fields in the PA Global Edit window

In the following field	Enter the following information
SMTP Server	SMTP server name
SMTP Port	SMTP port
Email from	The name that appears as the sender of proactive alerts
Email to	The email address of the recipient of proactive alerts You can add multiple recipients by using a comma to separate the addresses.
Snooze	Add the time (in minutes) for a pause between two consecutive proactive alerts. The snooze time that you configure is applicable for all components.

6. Complete the tasks in the following table to add a Canister server, data service, or PCA server for configuring proactive alerts:

Table 28. Adding a Canister, data service, or PCA

To add...	Complete these tasks in the Proactive Alerts Configuration window...
Canister	1. In the **Canister Server Rules** area, click **Add Canister**. The Add Canister window opens. 2. In the **Canister Servers** list, select a Canister. 3. In the **Port** field, enter the Canister port and click **OK**.

Table 28. Adding a Canister, data service, or PCA (continued)

To add...	Complete these tasks in the Proactive Alerts Configuration window...
Data service	1. In the **Data Services Rules** area, click **Add Data Service**. The Add Data Service window opens. 2. In the **Data Service Servers** list, select a data service. 3. In the **Port** field, enter the data service port and click **OK**.
PCA	1. In the **PCA Rules** area, click **Add PCA**. The Add PCA window opens. 2. In the **PCA Server** list, enter the host name, which is the fully qualified domain name (FQDN), of the PCA server. 3. In the **Port** field, enter the PCA port and click **OK**.

Adding alert conditions

Use the Proactive Alerts Configuration window to add proactive alert conditions for IBM Tealeaf components.

Complete the following steps to add proactive alert conditions:

1. On the TMS home page, go to the **Worldview** tab.
2. Expand your master server and then expand the **Tealeaf Management Server** node.
3. Click the **Proactive Alerts Configuration** *MasterServerName* link. The Proactive Alerts Configuration window opens.
4. Ensure that the **Turn on Proactive Alerts** check box is selected. If you want to turn off the Proactive Alerts feature, clear the **Turn on Proactive Alerts** check box.

   Important: If you have configured proactive alerts for any component, and if you clear the **Turn on Proactive Alerts** check box, proactive alerts are not generated for any component.
5. Complete the tasks in the following table to add an alert condition for a Canister, data service, or PCA:

   Note: An alert condition is an expression, and can have one of the following operators:

   <

   >

   =

Table 29. Adding an alert threshold

To add an alert for...	Complete these tasks in the Proactive Alerts Configuration window...
Canister	1. In the **Canister Server Rules** area, click **Edit** against a Canister server. The Canister Expression Edit window opens. 2. In the **Expression** field, select an expression. 3. In the **Operators** field, select an operator. 4. In the **Threshold** field, enter a value for the expression, and click **OK**.
Data service	1. In the **Data Service Rules** area, click **Edit** against a data service. The Data Service Expression Edit window opens. 2. In the **Expression** field, select an expression. 3. In the **Operators** field, select an operator. 4. In the **Threshold** field, enter a value for the expression, and click **OK**.
PCA	1. In the **PCA Rules** area, click **Edit** against a PCA server. The PCA Expression Edit window opens. 2. In the **Expressions** field, click **Add**. The PCA Expression Edit opens. 3. In the **Expression** field, select an expression. 4. In the **Operators** field, select an operator. 5. In the **Threshold** field, enter a value for the expression, and click **OK**. 6. If you want to receive an alert when an SSL is not decrypted, in the PCA Edit window, select the **Alert on SSL not getting decrypted** check box. 7. If you want to receive an alert for a failed connection, select the **Alert on aged connection** check box. 8. Verify the PCA port and click **OK**.

6. Click **Save**.
7. Restart the TMS server.

   Note: Configuration for proactive alerts is not effective till the TMS server is restarted.

Editing alert conditions

Use the Proactive Alerts Configuration window to edit proactive alert conditions for IBM Tealeaf components.

Complete the following steps to edit proactive alert conditions:

1. On the TMS home page, go to the **Worldview** tab.
2. Expand your master server and then expand the **Tealeaf Management Server** node.

3. Click the **Proactive Alerts Configuration** *MasterServerName* link. The Proactive Alerts Configuration window opens.
4. Ensure that the **Turn on Proactive Alerts** check box is selected. If you want to turn off the Proactive Alerts feature, clear the **Turn on Proactive Alerts** check box.

   Important: If you have configured proactive alerts for any component, and if you clear the **Turn on Proactive Alerts** check box, proactive alerts are not generated for any component.
5. Complete the tasks in the following table to edit an alert condition for a Canister, data service, or PCA:

   Note: An alert condition is an expression, and can have one of the following operators:

   <

   >

   =

Table 30. Adding an alert threshold

To edit an alert for...	Complete these tasks...
Canister	You can edit an expression from the Canister Edit window. Complete the following steps to edit an expression: 1. In the **Canister Server Rules** area, click **Edit** against a Canister server. The Canister Edit window opens. 2. In the **Canister Expression Edit** field, select an expression and click **Edit**. The Canister Expression Edit window opens. 3. In the **Operators** field, select an operator for the expression. 4. In the **Threshold** field, change the value of the expression, and click **OK**.
Data service	You can edit an expression from the Data Service Edit window. Complete the following steps to edit an expression: 1. In the **Data Service Rules** area, click **Edit** against a data service. The Data Service Edit window opens. 2. In the **Data Service Expression Edit** field, select an expression and click **Edit**. The Data Service Expression Edit window opens. 3. In the **Operators** field, select an operator for the expression. 4. In the **Threshold** field, change the value of the expression, and click **OK**.

Table 30. Adding an alert threshold (continued)

To edit an alert for...	Complete these tasks...
PCA	You can edit an expression from the PCA Edit window. Complete the following steps to edit an expression: 1. In the **PCA Rules** area, click **Edit** against a PCA server. The PCA Edit window opens. 2. In the **Expressions** field, select an expression and click **Edit**. The PCA Expression Edit window opens. 3. In the **Operators** field, select an operator for the expression. 4. In the **Threshold** field, change the value of the expression, and click **OK**.

6. Click **Save**.
7. Restart the TMS server.

   Note: Configuration for proactive alerts is not effective till the TMS server is restarted.

Proactive alerts for Canister

IBM Tealeaf administrators receive email alerts when one or more values of the parameters that you configured for a Canister server are less than or exceed the threshold.

Proactive alerts are available for the following thresholds:

FreeDiskRemaining(%)
The space that is available on the Canister disk, in percentage.

Note: Although a Canister might have more than one hard disks attached, an alert is generated for the disk space utilization of the hard disk that is configured in the **LOCAL_DIRECTORY** parameter in the `ctsrvr.cfg` file. The `ctsrvr.cfg` file exists in the Canister installation directory.

MemoryUtilization(%)
Canister server RAM utilization in percentage.

CPUUtilization(%)
Canister Server CPU utilization in percentage.

Note: An alert for the Canister server CPU utilization is generated for the physical CPU server, and not the Canister server process (`ctree.exe`).

SpoolTimeSecondsAllowed(Seconds)
The amount of time (in seconds) that is required for PCA hits to be spooled by the Canister. If the spool time reaches the value that you defined in the TMS, an email alert is generated.

PCA hits are spooled if they are not processed by the Canister.

Proactive alerts for data services

IBM Tealeaf administrators receive email alerts when one or more values of the parameters that you configured for a data service are less than or exceed the

threshold. Proactive alerts can be configured for the Dimension and Aggregate tables in the IBM Tealeaf database. The Dimension and Aggregate tables are generated dynamically.

Proactive alerts are available for the following thresholds:

DimensionSize(MB)
The size of Dimension tables in MB. This parameter is applicable for all Dimension tables. A threshold is not configured for individual Dimension tables.

AggregateSize(MB)
The size of Aggregate tables in MB. This parameter is applicable for all Aggregate tables. A threshold is not configured for individual Aggregate tables.

DimensionGrowth(%)
The daily percentage of the growth of data in Dimension tables. This parameter is applicable for all Dimension tables. A threshold is not configured for individual Dimension tables.

AggregateGrowth(%)
The daily percentage of the growth of data in Aggregate tables. This parameter is applicable for all Aggregate tables. A threshold is not configured for individual Aggregate tables.

Important: The data collector updates the statistics of database growth for dimensions and aggregates once in a day. The interval between two consecutive data server proactive alerts can be one day.

Proactive alerts for PCA

IBM Tealeaf administrators receive email alerts when one or more values of the parameters that you configured for Passive Capture Application (PCA) are less than or exceed the threshold.

The Tealeaf Management System (TMS) creates XML files for the current hit of the PCA statistics in the *tealeaf_installation_location*\System\PA folder. The PA folder is created when the TMS receives a hit. The TMS saves XML files of the previous three days in the *tealeaf_installation_location*\System\PA\Archive folder. XML files that are older than three days are deleted.

Note: The recommended frequency of sending statistics hit to PAStatsLogger is one minute.

Proactive alerts are available for the following PCA thresholds:

TotalHitsDropped
If the number of total hits dropped while data is sent to a processing server reaches the value that you defined in the Tealeaf Management System (TMS), an email alert is generated. The email alert contains a list of all the processing servers where dropped hits reach the threshold.

Aged Connections
If a connection with any processing server fails, an email alert is generated. The email alert contains a list of all the processing servers where connections fail.

ReassdHighCPUUtilization(%)
If the processor CPU utilization (in percent) of a reassd PCA process reaches the threshold value for three consecutive times, an alert is generated.

PipelinedHighCPUUtilization(%)
If the processor CPU utilization (in percent) of a pipelined PCA process reaches the threshold value for three consecutive times, an alert is generated.

SSL is not getting decrypted
If the secure sockets layer (SSL) traffic is not being decrypted by the PCA, an email alert is generated.

Patch Management with IBM Endpoint Manager

You can use IBM Endpoint Manager to run patches on your entire IBM Tealeaf network.

The following list summarizes the key capabilities of IBM Endpoint Manager:

- Automate patch management across a distributed setup.
- Apply the correct patch to the correct node.
- Achieve greater visibility in patch compliance with flexible, and almost real-time monitoring and reporting.
- Gain almost real-time visibility and control from a single management console.
- Reduce patch application cycles from weeks to hours.

You can obtain the IBM Endpoint Manager license key through an email that provides information about using the Flexera Entitlement Licensing for IBM Endpoint Manager. Flexera Entitlement Licensing is a software licensing, compliance, and entitlement management solution through which you can obtain the IBM Endpoint Manager license key.

Applying patches on the IBM Tealeaf network

Use IBM Endpoint Manager to run IBM Tealeaf patches on all IBM Tealeaf servers in your network, except the CX Passive Capture Application (CX PCA) server. You cannot run patches on CX PCA through IBM Endpoint Manager because CX PCA requires two-way communication with the server.

Complete the following tasks before you apply patches on your IBM Tealeaf network:

1. Install IBM Endpoint Manager version 9.1. For more information about installing IBM Tealeaf version 9.1, see the sections in the following table in *IBM Endpoint Manager Version 9.1 Installation Guide*.

Table 31. Sections for reference in IBM Endpoint Manager Version 9.1 Installation Guide

To do this...	**In** *IBM Endpoint Manager Version 9.1 Installation Guide***, see...**
Know prerequisites for installing IBM Endpoint Manager	"Assumptions and requirements"
Understand types of installation	"Types of installation"
Understand the activities to be completed before installation	"Before installing"

Table 31. Sections for reference in IBM Endpoint Manager Version 9.1 Installation Guide (continued)

To do this...	**In** *IBM Endpoint Manager Version 9.1 Installation Guide*, **see...**
Install on Windows systems	"Installing on Windows systems"
Install on Linux systems	"Installing on Linux systems"
Complete post-installation and configuration tasks	"Post-installation configuration steps"

2. Deploy IBM Endpoint Manager client version 9.1 on every computer where you want to apply IBM Tealeaf patches.

   Use the IBM Endpoint Manager Client Deploy tool to deploy IBM Endpoint Manager clients at IBM Tealeaf nodes. All the computers where you deploy IBM Tealeaf client are called agents, and become the clients of the IBM Tealeaf server. For more information, seethe sections in the following table in *IBM Endpoint Manager Version 9.1 Installation Guide*

Table 32. Sections for reference for client deployment in IBM Endpoint Manager Version 9.1 Installation Guide

To do this...	**In** *IBM Endpoint Manager Version 9.1 Installation Guide*, **see...**
Know more about installing the Client Deploy Tool	"Installing the Client Deploy Tool"
Know more about installing the clients	"Installing the clients"
Understand how to use the Client Deploy Tool	"Using the Client Deploy Tool"
Understand how to manually install clients	"Installing the Client Manually"

If you have dependency rules for applying patches, you must create an automation plan. The following points are examples of dependency rules in an IBM Tealeaf network:

- Start patch application on node B after the patch is successfully applied on node A. Else stop patch application.
- Start patch application on the server group A after the patch is successfully applied to all servers in server group B. Else stop patch application.

For more information about automation plans, see the *Endpoint Manager for Server Automation* guide for version 9.1.

Note: You must install the Automation Plan Engine before you can create an automation plan. For more information, see "Installing the Automation Plan Engine" in *Endpoint Manager for Server Automation*.

Keep in mind the following points:

- You must store all the patches in the following directory.

  `Installation folder\BES Server\wwwrootbes\Uploads`
- Make sure that you create a directory to download fixlets on your local server.

Complete the following steps to apply patches on all IBM Tealeaf servers:

1. Once you receive information about the availability of a patch, download the patch and the corresponding fixlet from IBM Fix Central (http://www-933.ibm.com/support/fixcentral/).

   Download the patch in the following folder:

   Installation folder\BES Server\wwwrootbes\Uploads

   Store the fixlet in the directory that you created for storing fixlets.
2. Open the IBM Endpoint Manager server console.
3. On the left pane, click and expand the **All contents** tab.
4. Click **Select Sites > Master Action Site > Fixlets and Tasks**.
5. Use the **Import** dialog to import a fixlet.

   You can open the **Import** dialog by clicking **File > Import**.
6. In the **Fixlets and tasks** pane, select the fixlet that you imported.

   Note: If necessary, you can edit the fixlet. For more information, see *IBM Endpoint Manager Version 9.1 Console Operator's Guide*.
7. In the **Fixlet:*Fixlet name*** pane, click the **Take Action** tab to open the Take Action window.
8. Select the **Select Devices** check box in the **Target** tab. Expand the **Applicable Computers** node to see a list of computers that are eligible for installing the patch.
9. From the right pane, select the nodes where the patch must be installed and click **OK**.

After clicking **OK**, patch execution starts. You can use the IBM Endpoint Manager console to monitor the progress of patch application. If the patch application succeeds, the status on the IBM Endpoint Manager console reflects "Fixed". If the patch application fails, the status reflects "Failed".

You can also view the status of patch application by using IBM Endpoint Manager Web Reports. The web reports contain details about all the patches that are applied on all IBM Tealeaf nodes in your network.

To view IBM Endpoint Manager Web Reports, go to **Start>IBM Endpoint Manager>BES Web Reports**. You can also view Web Reports at http://*servername*/webreports.

For more information on IBM Endpoint Manager Web Reports, see the "Web Reports" section in *IBM Endpoint Manager Version 9.1 Console Operator's Guide*.

Chapter 5. TMS Jobs Tab

This section describes how to use the **Jobs** tab to add, view, or cancel scheduled jobs for various TMS tasks.

Overview

The Jobs tab is used to add, view, or cancel scheduled jobs that perform various TMS tasks. All defined TMS jobs are listed on the right side, and the currently selected job is displayed on the left side.

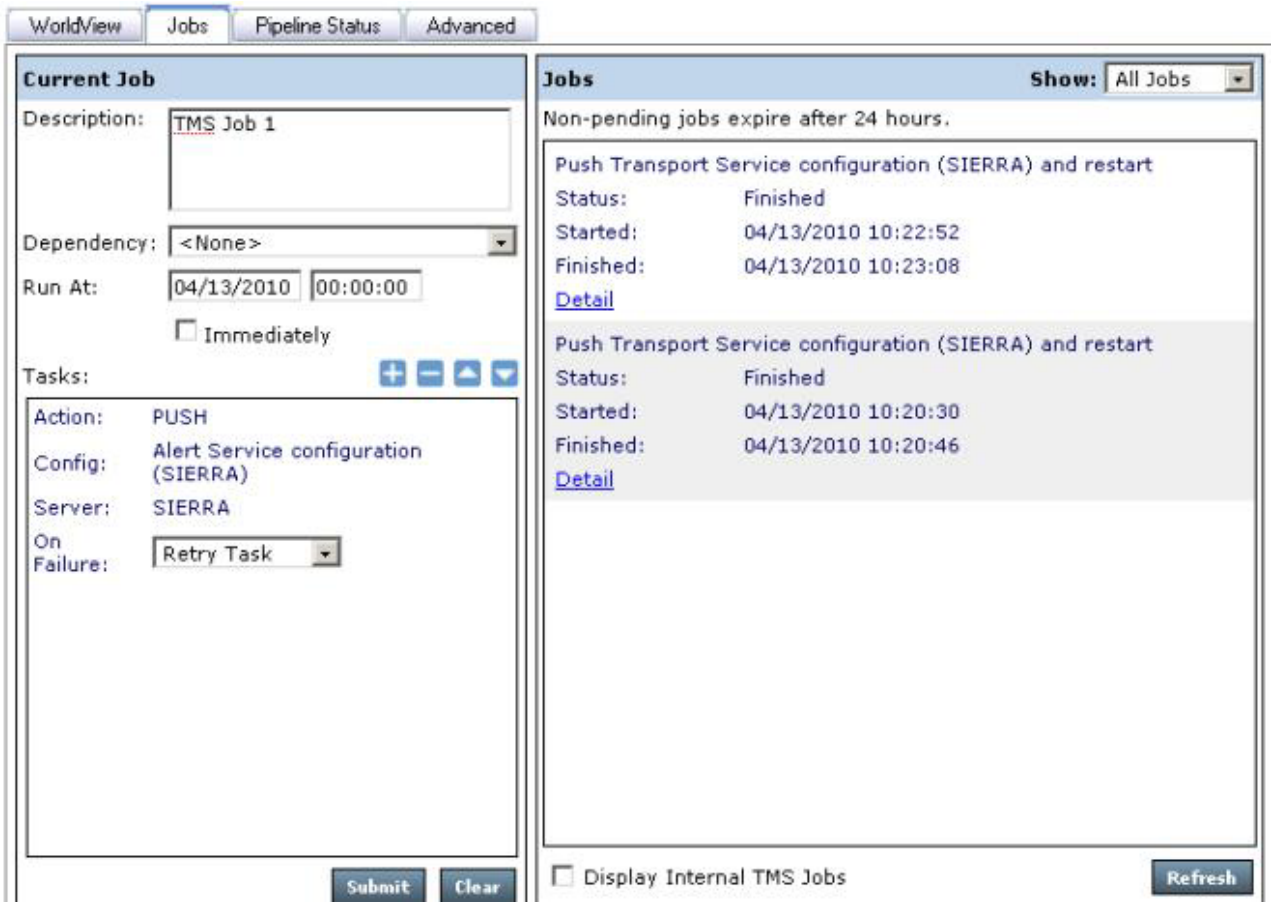

Figure 80. Jobs tab

On the left side, you can review the currently selected job. A TMS **job** is a sequence of tasks that TMS can perform on its master or slave servers. Tasks may include the following:

- Stop a server or service
- Start a server or service
- Cycle a server or service
- Assign a configuration
- Push a configuration to other server in the TMS environment

© Copyright IBM Corp. 1999, 2014

The fields in the Current Job pane are as follows:

Field Description

Description
This field is used to enter a description for the job. A meaningful description helps to later identify a job in the job list.

Dependency
In some cases, you may wish to execute a job after another job has completed successfully. The dependency field allows you to select a job that must complete before the current job is started.

Run At
This field allows you to specify a date and time at which the job will be run. If the date or time is in the past, then the job will be run immediately (subject to dependency limitations). By default, this field contains today's date with a time of `00:00:00`, which will cause it to run as soon as it's submitted.

Run Immediately
Click this checkbox to run the job as soon as you submit it. When this option is selected, you cannot specify a time at which to run it.

Tasks Pane

In the Tasks pane, you can create a new job or add tasks to an un-submitted job. Tasks are executed in the listed order.

- To add a task to the currently selected job, click the + button. See "Configuring tasks."
  - To create a new job, enter a new name in the Description and begin adding tasks. See "Configuring tasks."
- You can create dependencies between jobs, so that you can create sequences. When a job is selected from the Dependency drop-down, the currently specified job does not execute unless the dependency job successfully completed.
  - To create a sequence of jobs, configure them to execute at the same time, with each job dependent on the previous job in the sequence.
- To remove a task from the currently selected, job, select the task and click the **-** button.
- To remove all tasks from the current job, click **Clear** at the bottom of the Current Job pane.
- To change the order of execution of tasks in a job, select the task to move and then use the up and down arrows to re-order it. When you have finished editing the job, click **Submit**. The job is submitted to TMS for processing.
- To remove a job, select the scheduled or in-process job in the Jobs panel. Then, click the Cancel button.

Configuring tasks

In the Configure Task dialog, you can specify the type of task to add. The following table describes the available fields, some of which may not be shown depending on the selected Action:

Field Description

Action
The type of action to perform, including the following:

- Push Configuration
- Assign Configuration
- Start Component
- Stop Component
- Restart Component
- Archive History - See “Archive History” on page 294.

On Failure
Determines what happens if the action fails, which can include the following:

- Retry Task
- Stop Job
- Ignore Failure

Server
The server for which the action is to be taken.

Component
The component for which the action is to be taken.

Config
The configuration to be pushed (only shown for Push Configuration tasks).

Path The path to which the archived data is written (for Archive History tasks).

Age Cutoff (Days)
The number of days of history to keep (for Archive History tasks). History records older than the specified number of days are included in the archive.

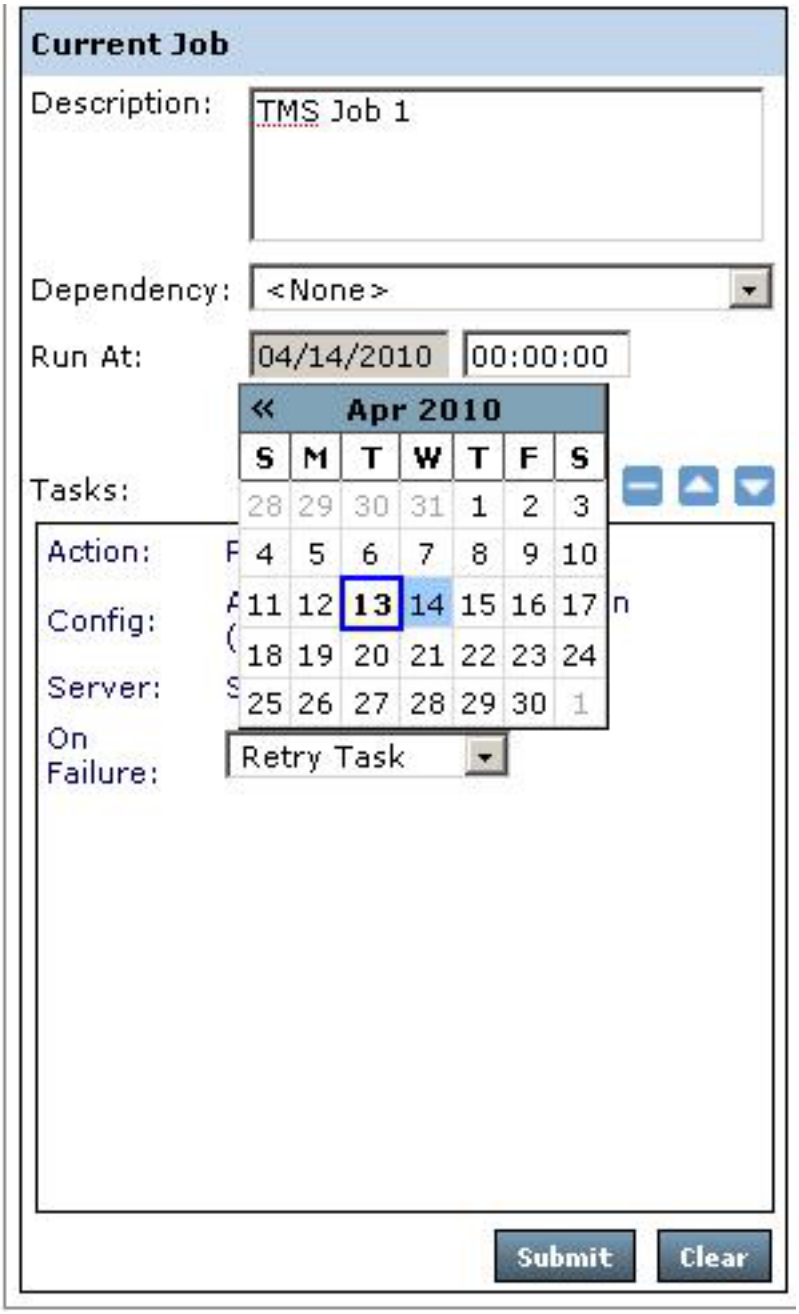

Figure 81. Current job

Archive History

For a specified date range, you can move old configuration versions and assignments to backup files. When files are successfully backed up, the source data is removed from the database tables to which TMS refers. This archiving removes outdated configuration information and reduces unnecessary bloat due to a large number of configuration changes.

Jobs Pane

On the right side of the Jobs tab, you can review all scheduled and completed jobs.

Note: After a job has executed, it is removed from the display 24 hours later.

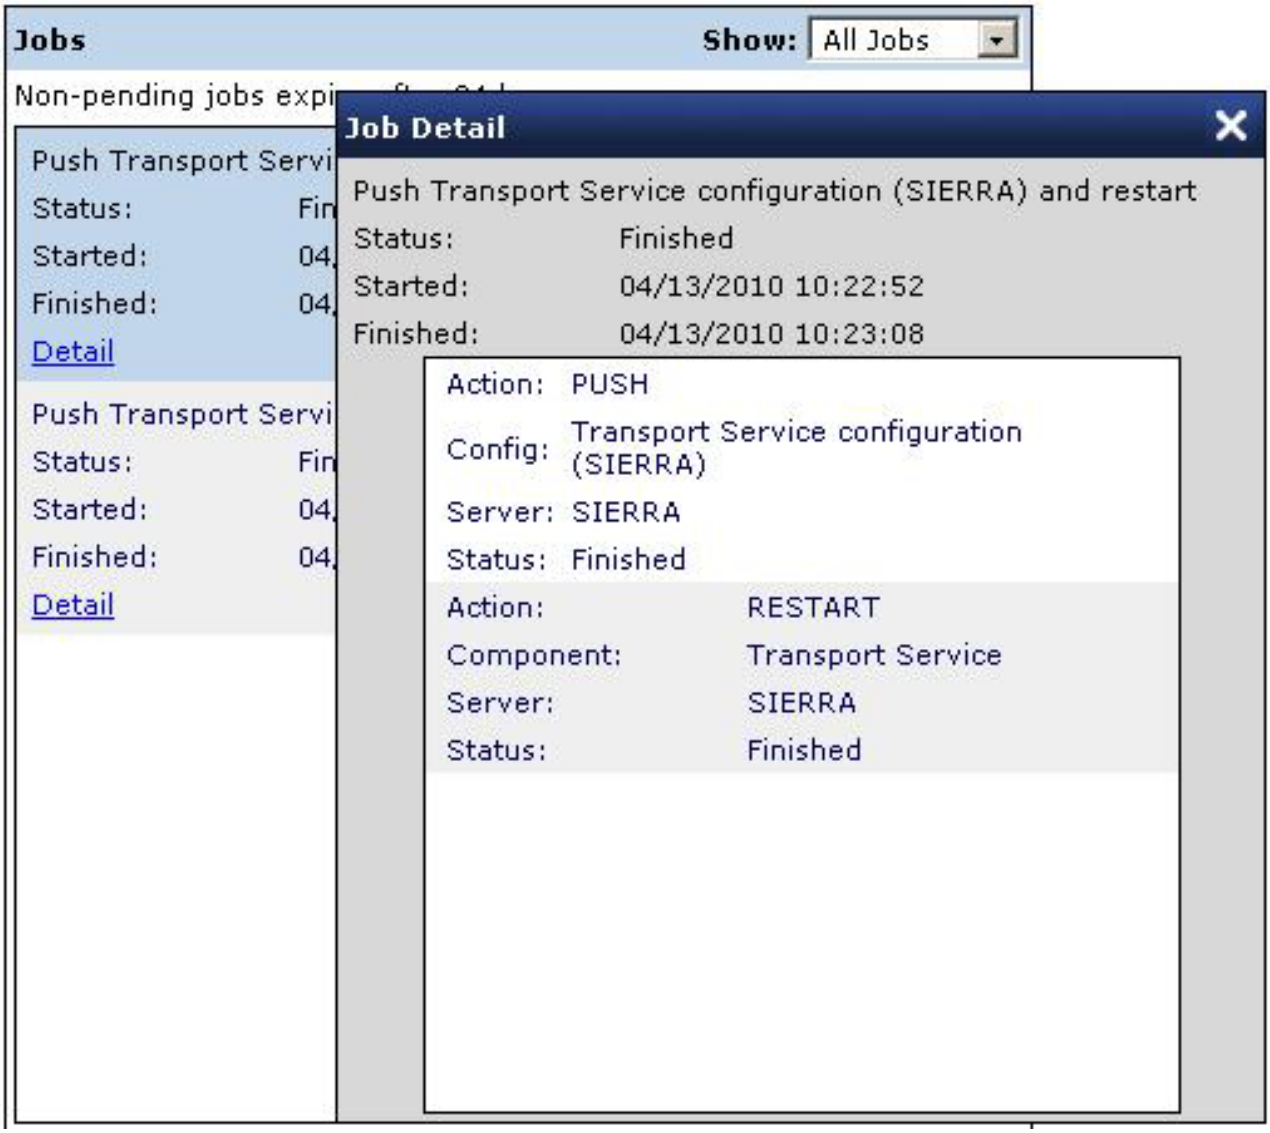

Figure 82. All Jobs details You can change the displayed jobs by selecting one of the following options from the Show drop-down:

- `Pending` - all jobs that have been defined and have not yet executed.
- `Processed` - all jobs that have already executed.
- `All Jobs` - all scheduled and all completed jobs.
- To toggle the display of jobs internal to TMS, click the Display Internal TMS Jobs checkbox.
- To view the component tasks of a job, click the Detail link.
- To cancel a scheduled or partially processed job, click the icon next to the job.
- To refresh the display in the Jobs pane, click **Refresh**.

Chapter 6. TMS Pipeline Status Tab

This chapter describes how to use the **Pipeline Status** tab to monitor Tealeaf capture and Transport Service pipelines.

Overview

Through the Pipeline Status tab you can monitor Tealeaf capture and Transport Service pipelines, including the processing by each agent in the pipeline, for all Windows-based Tealeaf servers in your environment. Throughput performance is also available for individual session agents and for the entire pipeline.

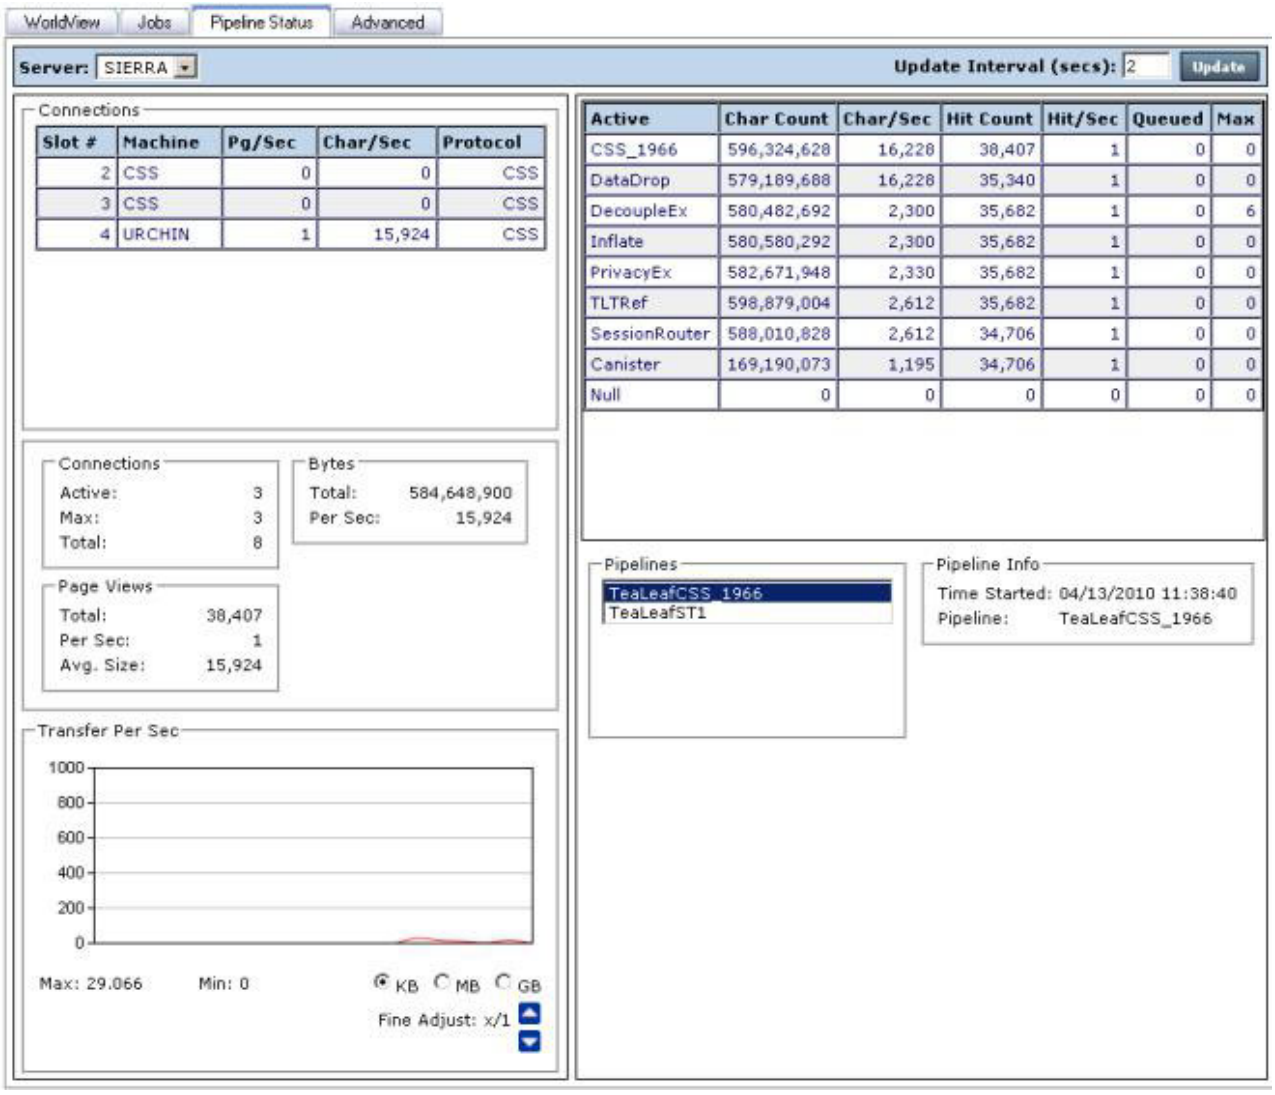
WorldView
Jobs
Pipeline Status
Advanced
Server: SIERRA
Update Interval (secs): 2
Update
Connections
Slot # | Machine | Pg/Sec | Char/Sec | Protocol
2 | CSS | 0 | 0 | CSS
3 | CSS | 0 | 0 | CSS
4 | URCHIN | 1 | 15,924 | CSS
Active | Char Count | Char/Sec | Hit Count | Hit/Sec | Queued | Max
CSS_1966 | 596,324,628 | 16,228 | 38,407 | 1 | 0 | 0
DataDrop | 579,189,688 | 16,228 | 35,340 | 1 | 0 | 0
DecoupleEx | 580,482,692 | 2,300 | 35,682 | 1 | 0 | 6
Inflate | 580,580,292 | 2,300 | 35,682 | 1 | 0 | 0
PrivacyEx | 582,671,948 | 2,330 | 35,682 | 1 | 0 | 0
TLTRef | 598,879,004 | 2,612 | 35,682 | 1 | 0 | 0
SessionRouter | 588,010,828 | 2,612 | 34,706 | 1 | 0 | 0
Canister | 169,190,073 | 1,195 | 34,706 | 1 | 0 | 0
Null | 0 | 0 | 0 | 0 | 0 | 0
Connections
Active: 3
Max: 3
Total: 8
Bytes
Total: 584,648,900
Per Sec: 15,924
Page Views
Total: 38,407
Per Sec: 1
Avg. Size: 15,924
Pipelines
TeaLeafCSS_1966
TeaLeafST1
Pipeline Info
Time Started: 04/13/2010 11:38:40
Pipeline: TeaLeafCSS_1966
Transfer Per Sec
1000
800
600
400
200
0
Max: 29.066
Min: 0
KB
MB
GB
Fine Adjust: x/1

Figure 83. TMS Pipeline Status tab

The **Servers** drop-down at the top of the window can be use to select the server to monitor.

© Copyright IBM Corp. 1999, 2014

Note: Only Windows-based servers with capture components or Transport Service are included in the Servers drop-down list.

By default, the selected server is polled for stats every two seconds via TMS.

- To change the update interval, enter a new positive figure in seconds in the Update Interval textbox. Then, click **Update**.

Connections

The Connections table shows active connections to the Transport Service, their throughput, the protocol in use, and the name of the machine from which the connection originates.

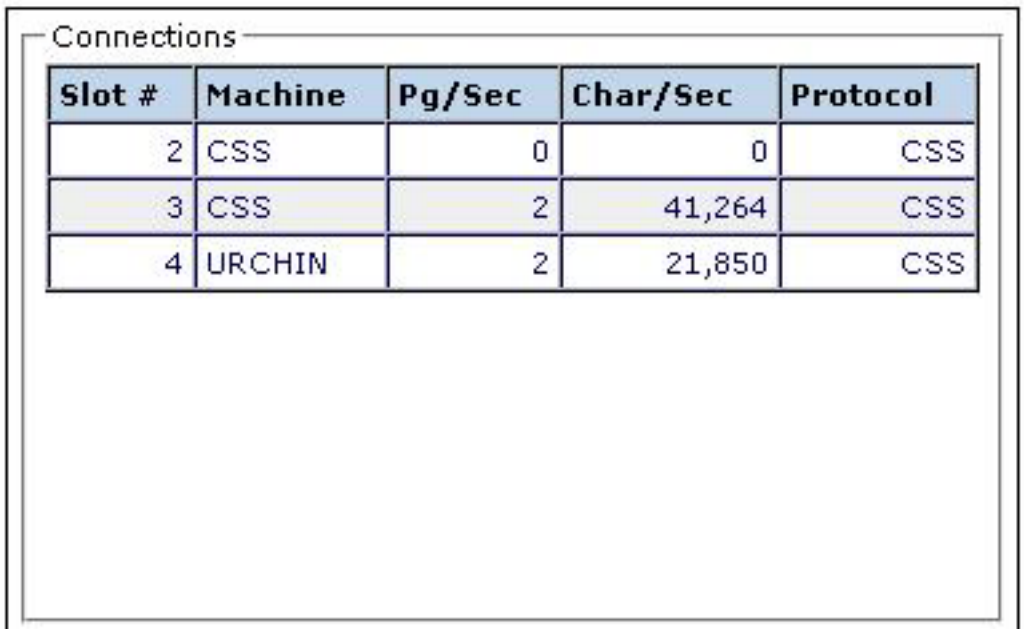

Figure 84. Connections Table

- All current connections are permanently displayed and cannot be added, removed or selected.
- The `slot#` column contains the memory address in Tealeaf's shared memory occupied by the connection.

Below the Connections table are statistics that reflect the overall state of all connections on the list.

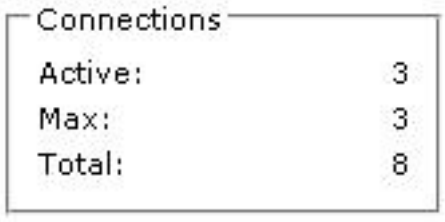

Figure 85. Connection Statistics

- **Active:** The number of active connections at any given time.

- **Max:** The maximum number of connections allowed.
- **Total:** The total number of connections seen since the application was launched.

Next to the Connection Stats, the Bytes area shows the total throughput of all connections:

Figure 86. Statistics on Bytes

- **Total:** The total number of bytes that have traveled through all connections.
- **Per Sec:** The number of bytes per second that are traveling through all connections.

Below the connections stats box, there is a set of stats regarding page transfers:

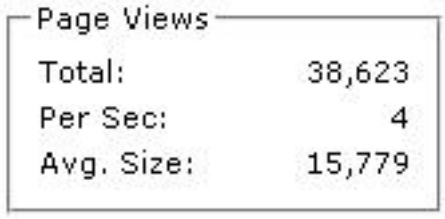

Figure 87. Page Statistics

- **Total:** The total number of pages seen since the launch of the utility.
- **Per Sec:** The number of pages seen in the last second.
- **Avg. Size:** The throughput over the last second divided by the number of pages seen in the last second.

Pipelines

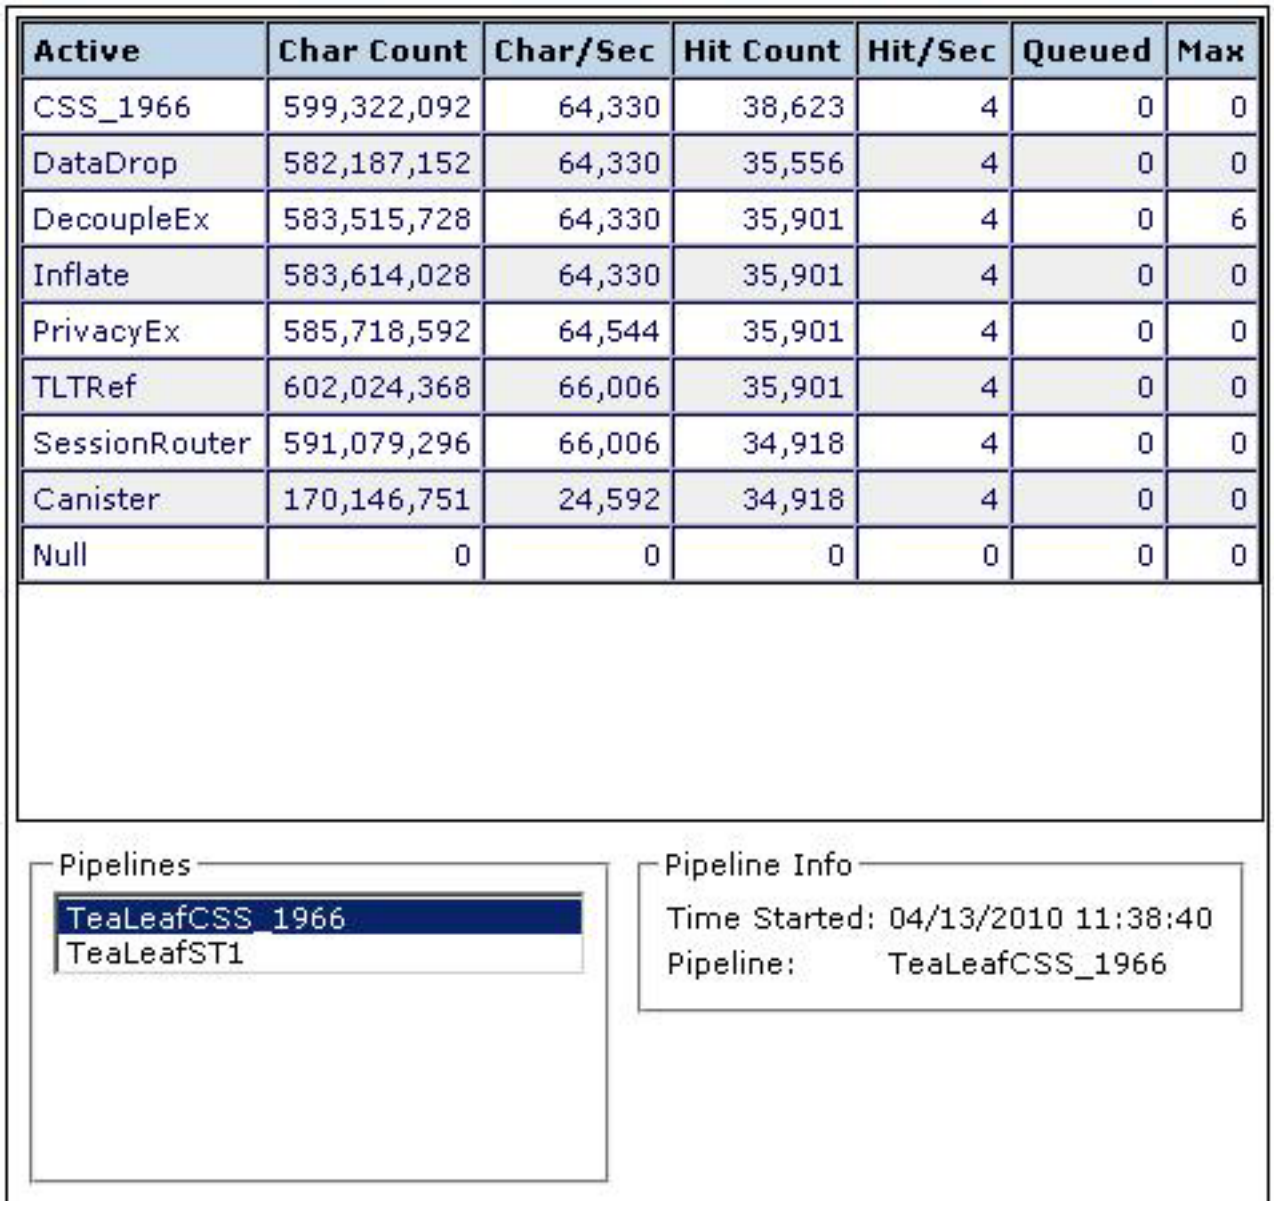

Active	Char Count	Char/Sec	Hit Count	Hit/Sec	Queued	Max
CSS_1966	599,322,092	64,330	38,623	4	0	0
DataDrop	582,187,152	64,330	35,556	4	0	0
DecoupleEx	583,515,728	64,330	35,901	4	0	6
Inflate	583,614,028	64,330	35,901	4	0	0
PrivacyEx	585,718,592	64,544	35,901	4	0	0
TLTRef	602,024,368	66,006	35,901	4	0	0
SessionRouter	591,079,296	66,006	34,918	4	0	0
Canister	170,146,751	24,592	34,918	4	0	0
Null	0	0	0	0	0	0

Figure 88. Agents and Pipelines

The Pipeline table shows the different session agents in the order in which they appear in the currently selected pipeline.

- You can select the pipeline to monitor in the pipeline list box below the table. When monitoring stats for a particular server, verify that you have selected the appropriate pipeline.
  - When you select a pipeline in the list, the list of agents and their data are updated to reflect the selected pipeline. The last time the pipeline was started is displayed in the Pipeline Info pane.
- Transfer rates for the current pipeline are displayed in the graph in the lower-left corner of the screen. See "Transfers per Second" on page 301.

In the Pipelines table, you can review each session agent in the pipeline in the order of execution. After the name of the session agent, the following statistics are displayed:

Statistics
Description

Agent Name of the session agent.

Char Count
Number of characters processed by the session agent since the pipeline was started.

Char/Sec
The number of characters processed per second.

Hit Count
Number of hits processed by the session agent since the pipeline was started.

Hit/Sec
The number of hits processed per second.

Queued
The number of hits that are currently queued in memory for processing.

Max The maximum number of hits that have been queued in memory for processing since the pipeline was started.

- Pipelines can be created and edited in the Pipeline Editor. See Chapter 8, "TMS Pipeline Editor," on page 307.

Transfers per Second

In the Transfers per Second graph in the lower left corner, you can monitor the traffic rate in the currently selected pipeline.

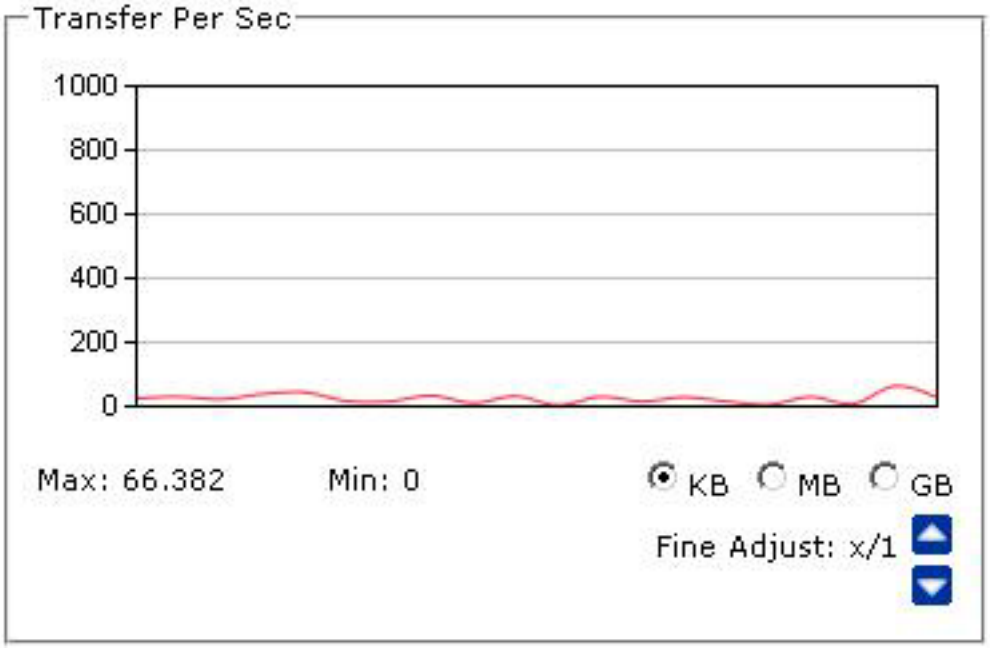

Figure 89. Pipeline Status graph

As data passes through the selected pipeline, you can review the traffic rate scrolling across the X-axis of the graph.

- Use the radio buttons to adjust the scale of the Y-axis to display kilobytes, megabytes, or gigabytes.
- Fine adjustments to the vertical scale can be made using the up and down arrows.

Chapter 7. TMS Advanced Tab

This section describes the advanced settings for TMS.

Overview

Through the Advanced tab, you can import and export TMS configurations and config summaries. Additionally, you can execute an auto-discovery operation, which scans TMS servers to discover new components and configurations that can be managed through TMS.

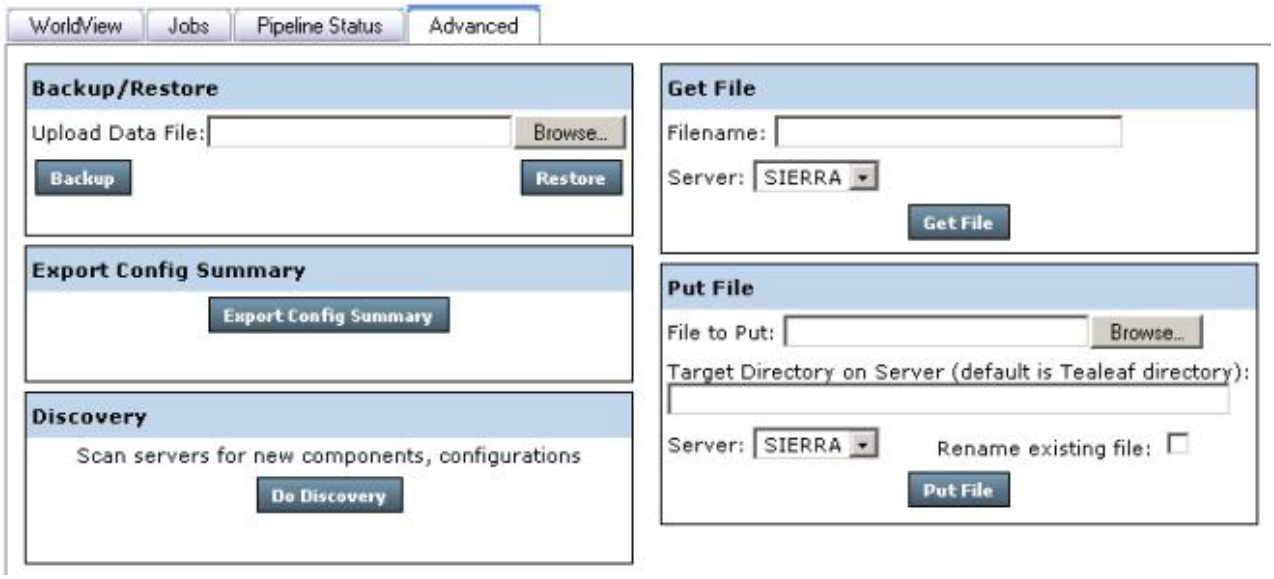

Figure 90. TMS - Advanced tab

Command
Description

"Backup/Restore" on page 304
Backup the TMS database to `.xml` format and restore backups as needed.

"Export TMS Config Summary" on page 304
Export a summary of the current Tealeaf configuration, as managed through TMS. Export is in `.xml` format.

"Backup/Restore" on page 304
Backup the TMS database to `.xml` format and restore backups as needed.

"Discovery" on page 305
Scan Tealeaf servers and directories for newly installed components, which TMS can manage.

"Get File" on page 305
Retrieve a specified file from the selected Tealeaf server for editing on your local computer.

"Put File" on page 305
Place a file that you have edited on your local computer into the specified directory on the selected Tealeaf server.

© Copyright IBM Corp. 1999, 2014

Backup/Restore

Through the TMS Advanced tab, you can perform full backups of the TMS datastore. If needed, backups can be restored to the TMS environment.

Backup TMS Database

Through the Advanced tab, you can create a backup of the entire TMS database to `.xml` format or restore from an earlier backup.

Note: All configuration data is encrypted in the TMS datastore and remains encrypted when written to a backup file.

To create a backup of the TMS datastore:

1. Click **Backup**.
2. Through the File Download dialog, you can specify where to save the backup file.

Restore TMS Database

To restore from a backup:

Note: Restoring a TMS backup overwrites all existing data in the TMS datastore and causes the TMS service to suspend requests temporarily until the restore process is complete.

Note: After the datastore has been restored on the TMS master, a command is sent to each TMS slave to re-synchronize its local data. Resynchronization may include pushing out a different version of one or more configurations. The components are not automatically restarted after this process, so it is usually a good idea to restart all Tealeaf services after a restore.

1. Enter the full path to the `.xml` file, or click the **Browse** button to navigate to the file.
2. To begin the restore, select the file and click **Restore**.

The TMS datastore can also be backed up via file copy. The TMS datastore files are located in the following directory:

```
<Tealeaf_Install_Directory>\System\TMSStore
```

Export TMS Config Summary

The Export Config Summary feature creates an `.xml` file containing all current components and configurations, which can be used to document the initial configuration of the Tealeaf software. This configuration file may also be requested by Tealeaf Customer Support to assist with troubleshooting.

To create a Config Summary Export file

1. Click **Export Config Summary**.
2. Click **Save** and browse to the directory where you wish to save the exported file.
3. The `.xml` file is saved in the selected location.

Discovery

TMS can be instructed to scan the Tealeaf environment for new components and configurations. When a discovery is executed, the Tealeaf Management System examines all registered Tealeaf servers and compares their configurations to the known TMS data, making additions, subtractions, and changes as needed.

To execute a discovery:

1. In the Advanced tab of TMS, click **Do Discovery**.
2. Select the servers to scan. Click **Go**.
3. When the scan is complete, TMS is updated.

File Management

In some enterprise environments, access to Tealeaf servers may not be available universally. A Tealeaf administrator can use the Get File and Put File commands to retrieve and write files to specific servers without having direct server access.

Note: Use of the Get File and Put File features must be in accordance with your enterprise procedures. For more information, please contact your IT department.

Get File

You can use the Get File command to acquire specified files from the selected server for local editing and resubmission.

To get a file

1. Enter the filename in the textbox.
2. Select the server from which to acquire the file.
3. Click **Get File**.
4. TMS scans the server's Tealeaf install directory and sub-directories for the file. You may choose to open the file or save it locally for editing.

Put File

After you have retrieved and edited a file locally, you can click the **Put File** command to post the file to the original server or to other servers in the Tealeaf environment.

Note: This command places the selected file in the destination directory of your choice and overwrites any version of the file that is already there. This feature should be used only by administrators familiar with the internal file structure of the Tealeaf install directory of the target system. For more information, please contact Tealeaf http://support.tealeaf.com.

To put a file:

1. To select the file, click **Browse...**. Navigate your local computing environment to select the file to put. Click **Open**.
2. Enter the target directory on the server where you wish to place the file.
   - To place the file in the Tealeaf install directory, leave a blank value in the textbox.
   - To place the file in the Reporting directory inside the Tealeaf install directory, enter:

     ```
     Reporting
     ```

3. Select the server on which to place the file.
4. To save the existing file as a renamed file, click the checkbox. The filename of the original is appended with `ORIG`.
5. Click **Put File**.
6. The selected file is placed on the selected server in the designated directory.
7. Repeat the above steps if you wish to propagate the file to other Tealeaf servers.

Chapter 8. TMS Pipeline Editor

This section describes how to use the TMS Pipeline Editor to create or edit the pipelines for the current server.

Overview

Through the Pipeline Editor, you can create or edit the pipelines for the current server.

1. To open the Pipeline Editor, select one of the following configurations:
   - Transport Service > Transport Service configuration
   - Canister > Event Bus configuration
2. When the item is selected, click **View/Edit**.
3. The configuration is opened in the Pipeline Editor.

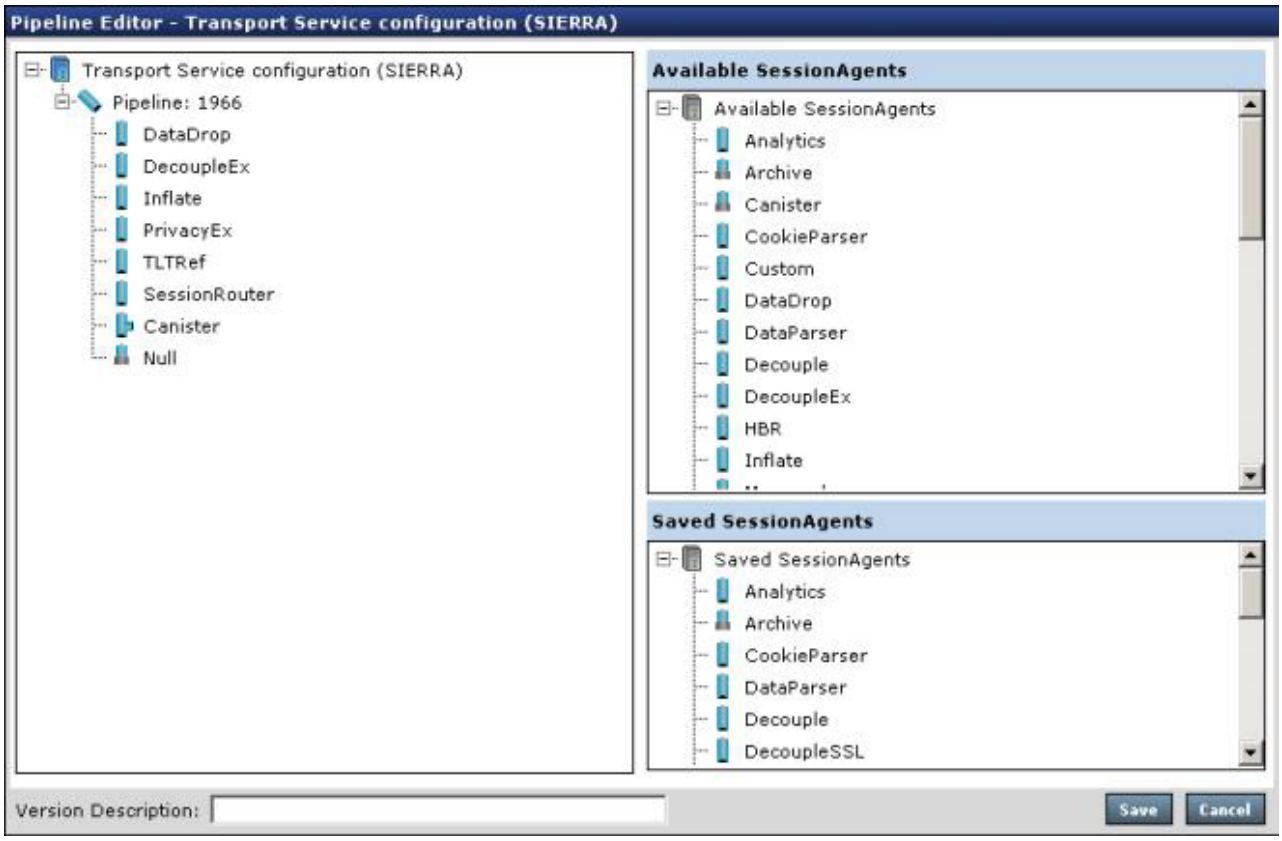

Figure 91. Pipeline Editor

In the Pipeline Editor, the pipelines defined for the current TMS server are displayed in the left panel. Beneath each pipeline, you can review the session agents in the pipeline. From top to bottom, they are displayed in the order of execution in the pipeline.

- You can add and remove session agents from pipelines by dragging and dropping them.

© Copyright IBM Corp. 1999, 2014

- **Service:**
  - "Editing Pipeline Globals"
- **Pipelines:**
  - "Adding Pipelines" on page 309
  - "Editing Pipelines" on page 310
  - "Removing Pipelines" on page 311
- **Session Agents:**
  - "Adding Session Agents" on page 311
  - "Editing Session Agents" on page 311
  - "Removing Session Agents" on page 312
  - "Saving Session Agents" on page 312

Editing Pipeline Globals

Through the Pipeline Editor, you can edit global settings that apply to all pipelines on the selected server.

To edit the pipeline globals, select the topmost node and click **Edit** at the bottom of the left pane.

In the dialog, you can edit the following settings.

Table 33. StatusMaster:

Setting	Description
`Application Name`	The name of the application. • For Tealeaf Event bus pipelines, set this value to `TealeafEventBus`. • For processing pipelines, set this value to `StatusMaster`.

Table 34. Globals for Transport Service Pipelines:

Setting	Description
`Maintenance Interval`	The frequency in seconds that a maintenance scan is performed.
`Max Data Size`	The maximum size in KB of any data accepted by the Tealeaf Capture Socket Service. Packets larger than this value are ignored. By default, this value is set to `2048` kilobytes. **Note:** This value should not be less than the values of any pipelines sending data, such as the IBM Tealeaf CX Passive Capture Application.
`Max Memory Size`	The maximum size in KB reserved on the server for the Transport Service. When the amount of memory used by the process exceeds this value, the service is automatically restarted. By default, this value is set to `1900` kilobytes. • The minimum value is `1024` MB (1 GB). • If this value is set to `0`, the minimum value is applied, and automatic restart of the service is disabled.
`Max Restart Count`	The number of times that the pipeline is permitted to restart within the number of seconds defined for `Restart Reset Interval`. • If the pipeline restarts more times than this value within one `Restart Reset Interval` period, the pipeline shuts down and must be restarted manually.

Table 34. Globals for Transport Service Pipelines: (continued)

Setting	Description
`Restart Interval`	The interval in seconds at which the pipeline checks to see if a restart is necessary.
`Restart Reset Interval`	The period of time in seconds over which the pipeline can attempt restarts. If the pipeline restarts more than `MaxRestart` times during this period, it is shut down.
`VM Free Disable Threshold`	The threshold percentage of free virtual memory below which the pipeline automatically shuts down. • To disable this check, set this value to `0`.
`VM Free Enable Threshold`	After a shutdown, the pipeline restarts when the percentage of free virtual memory exceeds this value. • This value should be greater than the value set for `VM Free Disable Threshold`, even if that value is set to `0`.

firewall

Firewall settings only apply to Transport Service pipelines. For more information on these settings, see "CX Pipeline Configuration" in the *IBM Tealeaf CX Configuration Manual*.

Additional documentation on all settings is available in the `TealeafCaptureSocket.cfg` file stored in the Tealeaf install directory on the server.

- See "CX Pipeline Configuration" in the *IBM Tealeaf CX Configuration Manual*.

Adding Pipelines

Note:

1. The Tealeaf Transport Service is a 32-bit application. Creating and using more than 20 pipelines in an individual Canister at one time may exceed the maximum available memory, which is up to 3GB. See “Editing Pipeline Globals” on page 308.
2. Depending on the objectives of the pipeline you are creating, there may be dependencies on the order in which the session agents are listed. Some session agents can only terminate a pipeline. For more general information on building a pipeline, see "CX Pipeline Configuration" in the *IBM Tealeaf CX Configuration Manual*.

1. To create a new pipeline, click the topmost node in the left panel.
2. Select **Add Pipeline**.
3. Enter a port number and select whether to use SSL in the pipeline.
4. Click **Apply**.
5. The pipeline is displayed in the left panel.
   - To edit the pipeline, select it and click **Edit**. See “Editing Pipelines” on page 310.
   - For more information on removal, see “Removing Pipelines” on page 311.
6. After the pipeline has been added and edited, you can begin adding session agents to it.
   - See “Adding Session Agents” on page 311.
   - See “Editing Session Agents” on page 311.

- See “Removing Session Agents” on page 312.

7. If your pipeline is receiving data, you can iterate on the configuration through TMS to tune performance. See “Pipeline Development Workflow” on page 313.

Building Pipelines

Through the Pipeline Editor, you can set up individual pipelines for each port to which the IBM Tealeaf CX Passive Capture Application is forwarding captured data.

- For more information on configuring the IBM Tealeaf CX Passive Capture Application to send traffic to a specific port, see "Configuration" in the *IBM Tealeaf Passive Capture Application Manual*.

By default, a single pipeline is created for you monitoring port 1966 using standard HTTP traffic. A set of pipeline session agents is added to the pipeline based on the installed components in your Tealeaf solution.

To build the pipeline:

- You can change firewall and other global settings for all of your pipelines. See “Editing Pipeline Globals” on page 308.
- You can change the port number and the use of standard or SSL traffic for any pipeline. See “Editing Pipelines.”
- You can continue to build the session agents of your pipeline. For a list of available session agents, see "Overview of the Capture Pipeline and Session Agents" in the *IBM Tealeaf CX Configuration Manual*.
- In this manner, you can build your pipeline until you reach one of the terminal pipeline session agents:
  1. The Canister session agent forwards the session data to a Tealeaf Short Term Canister. See "Canister Session Agent" in the *IBM Tealeaf CX Configuration Manual*.
  2. The Archive session agent forwards the session data to a Tealeaf Archive. See "Archive Session Agent" in the *IBM Tealeaf CX Configuration Manual*.
  3. The Null session agent discards the data. For debugging purposes.
  4. You can also point the end of a pipeline to the start of another one.
- Additional details may be managed through manual editing of the appropriate configuration file. See "CX Pipeline Configuration" in the *IBM Tealeaf CX Configuration Manual*.

Editing Pipelines

To edit the settings for any pipeline, select the pipeline node in the left panel. Then, click **Edit**.

Setting Description

Port The port used by the pipeline.

Note:

1. This value must be unique among all pipeline ports. By default, Tealeaf uses port 1966 for non-SSL traffic and 1967 for SSL traffic.
2. An SSL pipeline and a non-SSL pipeline should not both reference the same downstream session agent, especially if it is a canister. This situation may results in multiple threads.

Use SSL

When set to `True`, the pipeline uses SSL for transmission.

Note: Use of SSL may impact pipeline throughput.

Removing Pipelines

Note: When a pipeline is removed, all of its session agents are discarded as well. To preserve the configurations for session agents in a pipeline, drag them to the Saved SessionAgents panel first. See "Saving Session Agents" on page 312.

To remove a created pipeline, select it. Then, click **Remove Pipeline** at the bottom of the screen. Click **OK**.

Adding Session Agents

Session agents can be dragged from the Available Session Agents panel to the pipeline name to add them.

Session Agent icons in the Pipeline Editor:

- - A pipeline
- - A non-terminating session agent that passes data to another session agent downstream
- - A terminating session agent that is not the end of the pipeline
- - A terminating session agent

1. Drag the session agent over pipeline node and drop it.

   Note: You may be required to rename the session agent if an instance of the same session agent is already in the Saved SessionAgents pane.
2. Edit the settings for the session agent and click **OK**.
   - See "Editing Session Agents."
3. After multiple session agents have been added, you can reorder them within the pipeline or move them between pipelines using drag and drop.
   - For a list of available session agents and links to documentation, see "Available Session Agents" on page 312.

Editing Session Agents

To edit a session agent, select it in one of your pipelines. Then, click **Edit** at the bottom of the window. The session agent properties are displayed.

- Session agents are automatically named and numbered when they are added to a pipeline. To rename a session agent, double-click its name.
- To rename a session agent, click its node in a pipeline. Then, click **Rename**.

In the Edit SessionAgent dialog, you can change the properties for the selected session agent.

- The displayed properties depend on the selected session agent. For a list of available session agents and links to documentation, see "Available Session Agents."
- To change a setting, select it. A dialog displays the proper name, a brief description, and a method for changing its value. Select or enter a new value and click **Apply**.
- To add a custom configuration item for the session agent, click **Add Config Items**. See "Adding configuration items to session agents."

Adding configuration items to session agents

In the Add Config Item dialog, you can select configuration items that are available for the session agent to add to the current configuration.

- In the configuration file, these items are commented out. Selecting them enables them.
- To add a custom configuration item, click **Add Custom Config Item**. Enter a name for the custom configuration item and click **OK**. The item appears in the Add Config Item dialog.

  Note: Adding a custom configuration item only inserts the item into the configuration data. Customization of the affected `.DLL` is required. No data validation is performed on customer configuration items through TMS.
- After you have selected the item to add, click **OK**. The item is available for editing in the configuration.

Removing Session Agents

To remove a session agent, drag it from the pipeline.

- To retain it for later use, drag it to the Saved SessionAgents panel. See "Saving Session Agents."
- To remove it, drag it to the Available SessionAgents panel.

Saving Session Agents

You can use the Saved SessionAgents panel to retain configured session agents that you would like to use at a later time.

- When a session agent is dragged from a pipeline and dropped in the Saved Session Agents panel, it is removed from the pipeline and added to the Saved palette.
- When a session agent is dragged from the Available Session Agents panel and dropped in the Saved Session Agents panel, it is added to the Saved palette and renamed if its name is already in use.
- To edit a Saved session agent, drag it to a pipeline and click **Edit**.

Available Session Agents

The following session agents are currently available for inclusion in Tealeaf pipelines. Please review the links below for details on implementing session agents in your pipeline.

- "Adding a Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "Archive Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "Attribute Indexing Session Agent" in the *IBM Tealeaf CX Configuration Manual*

- "Canister Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "Cookie Parser Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "Data Drop Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "Data Parser Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "Decouple Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "Extended Decoupler Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "Extended Privacy Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "Health-Based Routing (HBR) Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "Inflate Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "JSON Mobile Parser Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "Managed Code Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "Null Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "Privacy Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "Real-Time Monitoring and Alert (RTA) Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "Response Tags to Request Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "RTA Split Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "Sessioning Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "Session Router Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "Socket Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "Statistics Logger Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "Tealeaf Reference Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "Tealeaf Reference Session Agent - Legacy Mode" in the *IBM Tealeaf CX Configuration Manual*
- "Tealeaf Sessioning Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "TimeGrades Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "TLI Session Agent" in the *IBM Tealeaf CX Configuration Manual*
- "URL Decode Session Agent" in the *IBM Tealeaf CX Configuration Manual*

Identify Deprecated Session Agents

When a system is upgraded, the deprecated session agents are not removed from the available session agents. Deprecated session agents can be identified in TMS Pipeline Editor. The deprecated session agent name is appended with (DEPRECATED) to highlight these session agents.

Pipeline Development Workflow

If your pipeline is receiving data, you can iterate on it through TMS to monitor performance.

1. Make changes to the pipeline through the Pipeline Editor.
2. In the Jobs tab, schedule a job to push the changes immediately to all servers. See Chapter 5, “TMS Jobs Tab,” on page 291.
3. Check the Pipeline Status tab to see if data is flowing through all session agents in the pipeline. See Chapter 6, “TMS Pipeline Status Tab,” on page 297.
4. Review the results in the destination system:

- If the pipeline ends at a Canister session agent, you can search in the Tealeaf Portal for active sessions or recently completed sessions to review the captured and processed data. See "Searching Session Data" in the *IBM Tealeaf cxImpact User Manual*.
- For Event Bus pipelines, you must review results in the destination system.

5. Repeat these steps until you are satisfied with your pipeline configuration.

Configuring the TMS Server

If needed, you can make changes to the configuration for the master Tealeaf Management System server.

Note: These settings are defined during installation and rarely need to be changed.

To access the TMS Server configuration:

1. Login to the Portal as an admin user.
2. From the Portal menu, select **Tealeaf** > **TMS**.
3. Expand the Tealeaf Management Server node.
4. Click **Tealeaf Management Server configuration**.
5. In the Config Actions panel, click **View/Edit**.
6. The following settings are available for review and editing:

   Note: If you modify the TMS server configuration and then submit a job to restart the service, you may see an error message in the Portal while the TMS service restarts. After TMS completes its restart, refresh the page in your browser to continue.

   Setting Description

   `Days to Keep Notifications`
   Number of days to retain notifications.

   `Log Level`
   Setting configures the desired level of logging verbosity for TMS.

   Note: Use log level `Debug` only for debugging issues related to TMS.

   `Master Port`
   Port number for master TMS server. Default value is `20000`.

   `Master TMS Server`
   Server hosting the master instance of TMS.

   Note: Tealeaf recommends that the TMS master server be set to the same server that hosts the Portal application, which reduces communication overhead and other issues. If you are installing on the Portal Server, please set this value to `localhost`.

   Note: When specifying the TMS master server, use a DNS-resolvable name that all Tealeaf servers can resolve. Do not use a fully qualified name.

Restarting TMS

The TMS server only needs to be restarted in the following situations:

- A configuration change is made to the Tealeaf Management Server. See "Configuring the TMS Server" on page 314 above.
- TMS has been upgraded.
- The TMS data dictionaries have been upgraded. Upgrades to TMS or its data dictionaries are uncommon.

Note: When TMS is restarted, Tealeaf users currently logged into the Portal may see an error message displayed in their browsers. This is normal. Refreshing the page clears the error.

1. Login to the Portal as an administrator.
2. To access TMS, select **Tealeaf** > **TMS** from the Portal menu.
3. In the TMS WorldView tab, select Servers view from the drop-down.
4. Select **Tealeaf Management Server**.
5. To restart the TMS server, click **Restart**.

Reinitializing TMS Datastore

If directed by Customer Support, you can re-initialize the Tealeaf Management System's datastore, which contains all configuration information for the components known to TMS.

Note: This step removes the Tealeaf Management System's datastore. If you have unsaved changes that you wish to apply, please save them and apply to all affected servers before you complete these steps.

1. If you have not done so already, you should identify the TMS Master Server. See "Configuring the TMS Server" on page 314.
2. Login to the server that is hosting the TMS Master Server as an administrator.
3. Open the Windows Services Control Panel.
4. Stop the Tealeaf Management Server service.
5. Navigate to the following directory:
   `<Tealeaf_install_directory>\System\TMSStore`
6. Delete all files in this directory.
7. In the Windows Services Control Panel, restart the Tealeaf Management Server service.
8. When the service is restarted, the `TMSStore` directory is repopulated via TMS discovery, and TMS is available through the Portal.

Troubleshooting

- If you are experiencing difficulties with TMS, please see "Troubleshooting - Tealeaf Management System" in the *IBM Tealeaf Troubleshooting Guide*.
- For more information on TMS configuration options, see the TMS Reference appendix.

Tutorial - TMS

This section contains How To tutorials on various aspects of the Tealeaf Management System.

- For more information on TMS in general, see Chapter 4, "Tealeaf Management System," on page 259.

TMS Components

The components Tealeaf Management system enables the following basic functionality:

Component
Description

WorldView tab
Select the server and the type of TMS-controlled objects on the server to view. Make changes to servers, configurations, and configuration items as needed.
- See "TMS WorldView Tab" on page 262.

Pipeline Editor
Accessed from within TMS WorldView tab, the TMS Pipeline Editor enables the creation and configuration of the Windows pipelines and their session agents, which process data captured by Tealeaf. See Chapter 8, "TMS Pipeline Editor," on page 307.

Jobs tab
Configure and execute jobs to assign configurations to servers in the Tealeaf environment, and start and stop services as needed.
- See Chapter 5, "TMS Jobs Tab," on page 291.

Pipeline Status tab
Monitor the performance of individual pipelines and session agents within each pipeline.
- See Chapter 6, "TMS Pipeline Status Tab," on page 297.

Advanced tab
Backup, restore, import and export of TMS configuration data.
- See Chapter 7, "TMS Advanced Tab," on page 303.

Install TMS

1. Specify TMS Master.
   - Use `localhost` or the machine name on master.
     - To acquire hostname, run `hostname` in console.
   - Use a DNS-resolvable name, preferably the machine name, of TMS master on slaves.
2. Install and start TMS on TMS master first, if possible.
3. Check event log `TLMgmtSrv*.log` for errors.
4. Configuration files picked up by TMS are copied to `<filename>.ORIG`.
   - `.ORIG` file is overwritten if it exists.
5. TMS datastore files are in `Tealeaf\System\TMSStore`.
   - TMS Master has full data store.
   - Slave TMS servers keep a subset of the data store in the same local location.
   - See "TMS WorldView Tab" on page 262.

Servers

This section describes how to disable, enable, group, and filter your servers.

Disable a Server

1. Select server to disable in the WorldView tab.

   Note: The TMS Master cannot be disabled.
2. Click **Disable Server** in the Server Actions area.
3. Click **Yes** to confirm.
   - See "TMS WorldView Tab" on page 262.

Enable a Server

1. Verify that **Show Disabled Servers** is selected at the bottom of the WorldView tab.
2. Select the server to enable.
3. Click **Enable Server** in the Server Actions area.
   - See "TMS WorldView Tab" on page 262.

Group a Set of Servers

Servers can be grouped by adding a common keyword to each server in a group.

1. Select a server to add to a group.
2. Click **Keywords...** in the Server Actions area.
3. Click **Add Keyword** in the Keywords dialog.
4. Select an existing keyword from the drop-down list or type a new keyword to use to identify the group. Then, click **OK**.
5. Add additional keywords if desired. Then, click **OK** to exit the Keywords dialog.
6. Repeat these steps for each server in the group.
   - See "TMS WorldView Tab" on page 262.

Filter the WorldView to Show a Group of Servers

This filtering uses keywords added to a group of servers to filter the servers shown in the WorldView.

1. Click the Filter checkbox at the bottom of the WorldView tab.
2. Click the Enable Filter checkbox if it is not already selected.
3. Click the Servers checkbox.
4. Change the operator dropdown to either match any or match all.
5. Click **Add** to add a keyword to the list.
6. Select the keyword for the server group from the dropdown list and click **OK**.
7. Click **OK** to apply the filter.
   - See "TMS WorldView Tab" on page 262.

Components

This section describes how to manage components.

Check a Component's Current State

1. Expand the node for the desired server in the WorldView tab.
2. A green icon next to the component name indicates the component is running.
   - A grey icon means that the component is stopped or not running.

3. Selecting the component displays the Component Info pane, which includes the current State of the component.
4. Click **Refresh** to refresh the component states for all servers.
   - See "TMS WorldView Tab" on page 262.

Start/Stop/Restart a Component

1. Expand the node for the desired server in the WorldView tab.
2. Select the component node you wish to control. The possible actions are displayed in the Component Actions section.
3. Click the appropriate button in the Component Actions section to perform the desired action immediately. Note that only components that are services or use identified services can be controlled.
4. You can optionally start, stop or restart a component as part of a job. See Chapter 5, "TMS Jobs Tab," on page 291.
   - See "TMS WorldView Tab" on page 262.

Add a New Component

Components are automatically added to TMS once they are discovered by the TMS service on the affected server.

1. Install the component on the desired server using the appropriate installer.
2. Click **Do Discovery** on the Advanced tab and select the affected server, or restart the TMS service on that server. Discovery is done automatically when the TMS service starts.
3. To refresh the WorldView tab, click **Refresh**.
   - See Chapter 7, "TMS Advanced Tab," on page 303.

Remove a Component

Components are automatically removed from TMS for the affected server if they are no longer detected during discovery.

1. Uninstall the component from the desired server using the appropriate installer.
2. Click **Do Discovery** on the Advanced tab and select the affected server, or restart the TMS service on that server. Discovery is done automatically when the TMS service starts.
3. To refresh the WorldView tab, click **Refresh**.
   - See Chapter 7, "TMS Advanced Tab," on page 303.

Group a Set of Components

Components can be grouped by adding a common keyword to each component in a group.

1. Select a component to add to a group.
2. Click **Keywords...** in the Component Actions area.
3. Click **Add Keyword** in the Keywords dialog.
4. Select an existing keyword from the drop-down list or type a new keyword to use to identify the group. Then, click **OK**.
5. Add additional keywords if desired, and then click **OK** to exit the Keywords dialog.
6. Repeat these steps for each component in the group.
   - See "TMS WorldView Tab" on page 262.

Filter the WorldView to Show a Group of Components

This filtering uses keywords added to a group of components to filter the components shown in the WorldView.

1. Click the Filter checkbox at the bottom of the WorldView tab.
2. Click the Enable Filter checkbox if it is not already selected.
3. Click the Components checkbox.
4. Change the operator dropdown to either match any or match all.
5. Click **Add** to add a keyword to the list.
6. Select the keyword for the component group from the dropdown list and click **OK**.
7. Click **OK** to apply the filter.
   - See "TMS WorldView Tab" on page 262.

Configurations

This section describes how to edit, revert, share, and copy configurations in the Pipeline Editor.

Edit a Configuration

1. Select the configuration to edit in the WorldView tab.
2. Click **View/Edit** (or **View/Edit Raw** to edit in raw mode) in the Config Actions area.
3. Make desired changes to the configuration.
4. Enter a version description in the field at the bottom of the configuration editor and click **Save** to save the updated configuration, or click **Cancel** to discard changes.
   - See "TMS WorldView Tab" on page 262.
5. If you saved changed to the configuration, select whether to add Push Config and/or Restart tasks.
6. Click **Add To Current Job** to add the selected tasks to a Job, which must be manually submitted, or **Add Tasks and Submit** to automatically submit a job with the selected tasks.
   - See Chapter 5, "TMS Jobs Tab," on page 291.

Revert a Configuration to a Previous Version

1. Select the desired configuration in the WorldView tab.
2. Select the version to revert to in the Versions dropdown list in the Config Info area.
3. Click **View/Edit** in the Config Actions area.
4. Enter a version description for the reverted configuration and click **Save**. It is saved as a new version.
   - See "TMS WorldView Tab" on page 262.
5. If you saved changed to the configuration, select whether to add Push Config and/or Restart tasks.
6. Click **Add To Current Job** to add the selected tasks to a Job, which must be manually submitted, or **Add Tasks and Submit** to automatically submit a job with the selected tasks.
   - See Chapter 5, "TMS Jobs Tab," on page 291.

Share a Configuration with Multiple Servers

1. Select the configuration to share in the WorldView tab.
2. Click **Assign (share)** in the Config Actions area.
3. If the configuration is not already used by all servers, then a list of servers not currently using the configuration is displayed.
4. Check the box next to the servers to assign the configuration to and click **Assign**.
   - See "TMS WorldView Tab" on page 262.
5. If the configuration was assigned to additional servers, select whether to add Push Config and/or Restart tasks.
6. Click **Add To Current Job** to add the selected tasks to a Job, which must be manually submitted, or **Add Tasks and Submit** to automatically submit a job with the selected tasks.
   - See Chapter 5, "TMS Jobs Tab," on page 291.

Make a Copy of a Configuration

1. Select the configuration to copy in the WorldView tab.
2. Click **Copy** in the Config Actions area.
3. Change the description to differentiate it from original.
4. Enter a version description and click **Copy**.

   Note: The new configuration can be found by switching to the Configs view using dropdown at the top of the WorldView pane and expanding the component to which the copied configuration belongs.
   - See "TMS WorldView Tab" on page 262.

Change the Description for a Configuration

1. Select the desired configuration in the WorldView tab.
2. Click **Edit Description** in the Config Actions area.
3. Edit the description and click **Save** to apply the change.
   - See "TMS WorldView Tab" on page 262.

Group a Set of Configurations

Configurations can be grouped by adding a common keyword to each configuration in a group.

1. Select a configuration to add to a group.
2. Click **Keywords...** in the Config Actions area.
3. Click **Add Keyword** in the Keywords dialog.
4. Select an existing keyword from the drop-down list or type a new keyword to use to identify the group. Then, click **OK**.
5. Add additional keywords if desired, and then click **OK** to exit the Keywords dialog.
6. Repeat these steps for each configuration in the group.
   - See "TMS WorldView Tab" on page 262.

Filter the WorldView to Show a Group of Configurations

This filtering uses keywords added to a group of configurations to filter the configurations shown in the WorldView.

1. Click the Filter checkbox at the bottom of the WorldView tab.
2. Click the Enable Filter checkbox if it is not already selected.
3. Click the Configs checkbox.
4. Change the operator dropdown to either match any or match all.
5. Click **Add** to add a keyword to the list.
6. Select the keyword for the configuration group from the dropdown list and click **OK**.
7. Click **OK** to apply the filter.
   - See "TMS WorldView Tab" on page 262.

Pipeline Editor

The Pipeline Editor is displayed whenever a pipeline configuration is edited. Pipeline Editor manages the Transport Server configuration (`TeaLeafCaptureSocket.cfg`), the Event Bus configuration (`TeaLeafEventBus.cfg`), as well as the various child pipeline configurations used by some of the Session Agents.

- The Tealeaf Event Bus can be used to deliver Tealeaf events to third-party systems for additional analysis and reporting. The Tealeaf Event Bus is a component of IBM Tealeaf cxConnect for Data Analysis, a separately licensable product for the IBM Tealeaf CX platform. See "Tealeaf Event Bus" in the *IBM Tealeaf cxConnect for Data Analysis Administration Manual*.

Edit Other Configuration Settings

1. Select the head node containing the configuration description and click **Edit** at the bottom of the editor, or double-click the head node. The Edit Config Sections dialog is displayed.
2. Click the + next to a section to expand it.
3. Click a configuration setting to change it. The Edit Config Item dialog is displayed. Change the value as desired and click **Apply** to keep the change.
4. When done making changes, click **OK** to keep the changes, or click the **Cancel** button to discard changes.
   - See "TMS WorldView Tab" on page 262.

Edit Pipeline Settings

This option is only available for configurations that can have multiple pipelines (e.g. `TeaLeafCaptureSocket.cfg`).

1. Select the desired pipeline node and click **Edit** at the bottom of the editor, or double-click the pipeline node. The Edit Pipeline Settings dialog is displayed.
2. Edit the settings as desired. Click **Apply** to keep the changes, or click **Cancel** to discard the changes.
   - See Chapter 8, "TMS Pipeline Editor," on page 307.

Edit SessionAgent Settings

1. Select the SessionAgent node and click **Edit** at the bottom of the editor, or double-click the SessionAgent node. The Edit SessionAgent dialog is displayed.
2. Click the + next to a group to expand it, if the SessionAgent has configuration item groups.
3. Click a configuration setting to change it. The Edit Config Item dialog is displayed. Change the value as desired and click **Apply** to keep the change.

4. When done making changes, click **OK** to keep the changes, or click the **Cancel** button to discard changes.
   - See Chapter 8, "TMS Pipeline Editor," on page 307.

Add a Pipeline

This option is only available for configurations that can have multiple pipelines (e.g. TeaLeafCaptureSocket.cfg).

1. Select the head node, the one with the configuration description.
2. Click **Add Pipeline** at the bottom of the editor. If there is no Add Pipeline button, then the current configuration doesn't support multiple top-level pipelines.
3. Enter the port number on which the pipeline listens.
4. Check the Use SSL checkbox to force incoming connections to use SSL.
5. Click **Apply** to add the new pipeline, or **Cancel** to discard it.
   - See Chapter 8, "TMS Pipeline Editor," on page 307.

Remove a Pipeline

This option is only available for configurations that can have multiple pipelines (e.g. `TeaLeafCaptureSocket.cfg`).

1. Select the pipeline to remove.
2. Click **Remove Pipeline** at the bottom of the editor.
3. Click **OK** to confirm the removal of the pipeline.

Note: There must be at least one pipeline in order to save the configuration.

See Chapter 8, "TMS Pipeline Editor," on page 307.

Add a SessionAgent to a Pipeline

1. Select a SessionAgent to add from either the Available SessionAgents list (SessionAgents with default settings) or the Saved SessionAgents list (SessionAgents with saved settings).
2. Drag the SessionAgent to the desired position in the pipeline tree. A line appears when you hover over the pipeline showing where the SessionAgent is to be placed.
3. After the SessionAgent is dropped in the pipeline, the Edit SessionAgent dialog is displayed, allowing you to change the settings for the SessionAgent.
4. Click **OK** to keep any changes made to the SessionAgent settings or **Cancel** to discard any changes made.
   - See Chapter 8, "TMS Pipeline Editor," on page 307.

Remove a SessionAgent from a Pipeline

1. Select the SessionAgent to remove.
2. Drag the SessionAgent to the Available SessionAgents list to remove the SessionAgent from the pipeline or to the Saved SessionAgents list to save SessionAgent instance for future use in this configuration.
   - See Chapter 8, "TMS Pipeline Editor," on page 307.

Move a SessionAgent within a Pipeline

1. Select the SessionAgent to move.

2. Drag the SessionAgent to the desired position in the pipeline tree. A line appears when you hover over the pipeline showing where the SessionAgent is to be placed.
   - See Chapter 8, “TMS Pipeline Editor,” on page 307.

Add a Custom SessionAgent

1. Select the appropriate base type in the Available SessionAgents list:
   - `Custom` - Custom native (C++) SessionAgent
   - `Managed` - C# or VB.NET SessionAgent hosted by SessionAgentCLR
   - `RulesEngine` - TCL script hosted by SessionAgentRulesEngine
2. Drag the SessionAgent to the desired position in the pipeline tree. A line appears when you hover over the pipeline showing where the SessionAgent is to be placed.
3. After the SessionAgent is dropped in the pipeline, the Edit SessionAgent dialog is displayed, allowing you to change the settings for the SessionAgent. Click **Add Config Items** to add custom configuration items for the SessionAgent.
4. Click **OK** to keep any changes made to the SessionAgent settings or **Cancel** to discard any changes made.
   - See Chapter 8, “TMS Pipeline Editor,” on page 307.

Jobs

This section describes how to manage and troubleshoot jobs.

Create a New Job

1. If there are exiting tasks in the Current Job pane, click **Submit** to submit the job or **Clear** to discard the current tasks.
2. Enter a description for the job in the Description fields.
3. Select a job on which to make this job a dependant, if desired. See “Make a Job Dependent on Another Job” on page 324.
4. Enter the date and time when you want the job to run, or if you want it to run as soon as possible, check the Immediately checkbox.
5. Add tasks to the job as described below.
   - See Chapter 5, “TMS Jobs Tab,” on page 291.

Add Tasks to a Job

1. Click the + button to display the Configure Task dialog.
2. Select the Action to take (other fields may change depending on the Action selected).
3. Select the On Failure action.
4. Select the Server that this job affects.
5. Select other application items (Component, Configuration, etc.).
6. Click **Save** to save the new task, or **Cancel** to discard it.
   - See Chapter 5, “TMS Jobs Tab,” on page 291.

Submit a Job

1. Verify that the description, scheduled time and tasks are what you want.
2. Click **Submit** to submit the job.

3. The submitted job should appear in the Jobs pane, unless the view is set to `Processed`.
   - See Chapter 5, “TMS Jobs Tab,” on page 291.

Cancel a Job

1. Verify that the Show option for the Jobs pane is set to `All Jobs` or `Pending`.
   - Processed jobs cannot be cancelled.
2. Click the icon in the upper right corner of the job in the Jobs pane to cancel it.
   - See Chapter 5, “TMS Jobs Tab,” on page 291.

View Job Status

Select the desired grouping in the Show option for the Jobs pane:

- `All Jobs` - Displays all jobs in the following order:
  - Completed jobs in reverse chronological order (most recently completed first)
  - Started jobs in reverse chronological order (most recently started first)
  - Scheduled jobs in chronological order by schedule date/time.
- `Pending` - Displays scheduled and incomplete jobs is the following order:
  - Started jobs in reverse chronological order (most recently started first)
  - Scheduled jobs in chronological order by schedule date/time.
- `Processed` - Displayed processed (and cancelled) jobs in reverse chronological order (most recently completed first)
  - Click the Detail link on the desired job to show the individual tasks for the job.

See Chapter 5, “TMS Jobs Tab,” on page 291.

Make a Job Dependent on Another Job

This option is used to verify that one job finishes processing successfully prior to a second job being executed. If job listed as a dependency fails, then any dependent jobs are not run. If a job is cancelled, all dependent jobs are also cancelled.

1. Before submitting the job, click the Dependency dropdown. All jobs that are eligible to be used as a dependency are displayed in the list.
2. Select a job from the Dependency dropdown.
   - See Chapter 5, “TMS Jobs Tab,” on page 291.

Troubleshoot a Failed Job

1. Locate the failed job in the Jobs pane. Verify that the Show option is set to `All Jobs` or `Processed`.
2. Click the Detail link for the failed job to show the task detail.
3. Find the task that failed and look for an error message in the Message field. This message should give some indication of what went wrong.
   - See Chapter 5, “TMS Jobs Tab,” on page 291.

Pipeline Status

This section describes how to view and troubleshoot the pipeline status.

View Pipeline Status for a Server

1. Select the Pipeline Status tab.
2. Select the desired server from the Server dropdown.
   - See Chapter 6, “TMS Pipeline Status Tab,” on page 297.

Switch to a Different Pipeline

Select the desired pipeline from the Pipelines list.

- See Chapter 6, “TMS Pipeline Status Tab,” on page 297.

Troubleshoot Pipeline Status

This section contains information to help troubleshoot issues that are related to the pipeline status.

Advanced

This section describes the troubleshooting options that are available under the **Advanced** tab when viewing the pipeline status.

Create a Backup of the TMS Data Store:

1. Select the Advanced tab.
2. Click the Backup button.
3. Click **Save** in the File Download dialog for your browser.
4. Select the location to save the backup file.
   - See Chapter 7, “TMS Advanced Tab,” on page 303.

Restore from a Backup of the TMS Data Store:

1. Select the Advanced tab.
2. Enter the full path and filename for the backup file in the Upload Data File field, or click **Browse** and navigate to the backup file. This file must be a file created using the Backup option.
3. Click **Restore**.
4. Click **OK** to confirm the restore or **Cancel** to quit.

   Note: During the restore process the TMS service is temporarily unavailable. Refresh the page in your browser to reestablish communication with TMS.
   - See Chapter 7, “TMS Advanced Tab,” on page 303.

Export a Configuration Summary:

1. Select the Advanced tab.
2. Click **Export Config Summary**.
3. Click **Save** in the File Download dialog for your browser.
4. Select the location to save the configuration summary file.
   - See Chapter 7, “TMS Advanced Tab,” on page 303.

Perform Discovery on a Server:

1. This option is used to detect added or removed components, or added configurations.
2. Select the Advanced tab.
3. Click **Do Discovery**.
4. Select the server(s) on which to perform the discovery process.

5. Click **Go** to initiate the discovery process.
6. Click **Refresh** on the WorldView tab to display any changes made during discovery.
   - See Chapter 7, “TMS Advanced Tab,” on page 303.

File Management: You can get files from any Tealeaf server in the environment for local editing and submission.

- For more information on getting files, see Chapter 7, “TMS Advanced Tab,” on page 303.
- For more information on pushing files to a server, see Chapter 7, “TMS Advanced Tab,” on page 303.

TMS Errors

This section describes TMS error messages that might be displayed and how to begin troubleshooting the errors.

Errors in Portal:
This section describes error messages that might be displayed in relation to the portal.

There was an error connecting to the TMS service:

- Verify that TMS service is running
- Check TMS access log for requests from the Portal.
  - See “Portal Logs” on page 227.

Error getting configuration data:
This message usually indicates a problem with the configuration data in TMS. Check other configurations. If all are showing this problem, then wipe out TMS datastore and start again.

Other servers are not appearing:

- Verify that TMS is started on other servers
- Verify that other servers have MasterServer set to actual TMS Master

Errors in Application Event Log:
This section describes error messages that might be displayed in the application event log.

<server> unable to connect to master <TMS master>:
Logged every 5 minutes while master can't be reached.

- Verify that TMS is running on the TMS Master Server.
  - TMS Master is configured through TMS and can be verified in TMS. See “Configuring the TMS Server” on page 314.
- In the Windows Registry, check the MasterServer value in:
  - 32-bit:

    ```
    HKLM\Software\TeaLeaf Technology\TLMgmtSrv
    ```
  - 64-bit:

    ```
    HKLM\Software\Wow6432Node\TeaLeaf Technology\TLMgmtSrv
    ```
- If the TMS Master Server is the same server hosting the Portal, set this value to `localhost`.

<server> error starting HTTPS server!:
Check rest of error message for information.

<server> error loading TMS configuration!:

- TMS configuration table is missing or corrupted.
- Reinitialize data store.
  - See "Configuring the TMS Server" on page 314.

<server> error loading data store!:

- Error while loading the data store.
- Check rest of error message for more information.
- Most likely, you must reinitialize data store. See "Configuring the TMS Server" on page 314.

<server> error building data store!:

- Error while building data store.
- Check rest of error message for more information.
- Verify that disk is not full.

<server> error initializing data store!:
Check rest of error message for more information.

<server> error reinitializing data store!:

- Check rest of error message for more information.
- This error may be encountered when reinitializing the data store from the command line.
  - You can manually reinitialize the data store. See "Configuring the TMS Server" on page 314.

<server> an exception was caught!:

- Check message for dump file location.
- Send dump file to Tealeaf Technical Support.

<server> error doing discovery of installed components and configurations!:

- Check rest of error message for more information.
- Most likely encountered on slave - error sending data to master.
- Most likely you must delete data store on slave, which then reloads from the master.

<server> error reinitializing data store!:
Check rest of error message for more information.

- See "Configuring the TMS Server" on page 314.

server ID not found:

- On TMS Master - either server was renamed or TMS Master server name was entered incorrectly.
- In Windows Registry, check MasterServer value in `HKLM\Software\TeaLeaf Technology\TLMgmtSrv` for correct server.
- Reinitialize data store using `TLMgmtSrv -reinit` command.

<server> TMS encryption key is different than key in TMS datastore. Please contact Tealeaf technical support.:

- May occur if TMS was uninstalled and then reinstalled without removing the old TMS data store (`Tealeaf\System\TMSStore` directory).

- Reinitialize data store to fix.
  - See “Configuring the TMS Server” on page 314.

Chapter 9. Advanced Administration Topics

This section provides additional information on the following administration topics. You should review these topics during initial configuration or upgrade of your Tealeaf system.

Authentication

For more information on authentication topics, see "TLTopic - Authentication" in the *IBM Tealeaf Topics*.

Supported Authentication Methods

The IBM Tealeaf CX Report Server Configuration tool contains four methods for authenticating user access. To access the configuration methods:

1. Login to the Tealeaf Portal as an administrator.
2. In the Portal menu, select **Tealeaf** > **TMS**.
3. Click the Tealeaf node.
4. Click **Shared configuration information**.
5. In the Config Actions, click **View/Edit**.
6. Click the Portal tab. The available methods are described below.
   - See "Configuring the Report Server" in the *IBM Tealeaf CX Configuration Manual*.

Portal (Database) Authentication

By default, the Portal is configured to use Portal authentication, which leverages the Portal's user management capabilities to control access. The Tealeaf Portal utilizes its own authentication mechanism, which is useful for organizations preferring to separate Tealeaf from their corporate authentication systems.

- **MD5** - This method only exists to support early customers configuring via the MD5 mechanism.

  Note: New Tealeaf systems should not deploy this authentication mechanism.
- **RC2** - Passwords are encrypted using RC2.
- **3DES** - Passwords are encrypted using 3DES (Default)
- **NT** - N/A

© Copyright IBM Corp. 1999, 2014

Figure 92. Portal Authentication login screen

NT (AD - Active Directory) Authentication

Under NT authentication mode, only valid NT accounts are permitted to log in to the Portal. Users submitting credentials that cannot be validated against the NT domain are not granted access.

- In Mixed Mode, NT users and accounts created through the Portal can access the Portal. See “Mixed Mode” on page 332.
- When a user performs a search, the Portal application passes the user's NT credentials to the Search Server, which validates those credentials against the NT domain. If those credentials are not valid, no data is returned from Search Server.

Note: Accounts that are registered as Administrator accounts in the NT domain are not granted administrator privileges automatically in the Portal. Those accounts must be assigned to the Admin group through the Portal.

Organizations that wish to integrate Tealeaf Portal logins with their existing AD groups may do so:

- **MD5/RC2/3DES** - When Allow Mixed Mode Logon is enabled, these three methods are available. Although the system is in NT mode, Mixed Mode allows a Portal authentication account to successfully login to the Portal.

  Note: With the exception of the designated user group type administrators (admin, revealadmin and viewadmin), Portal authentication accounts do not have access to event or session information, which are controlled by the search server under NT authentication. Mixed Mode allows an administrator to login and do account management regardless of the state of the NT authentication configuration.
- **NT** - This option should be used when NT/AD authentication is desired.

  Note:

  1. Active Directory user login names should not included any accented characters.
  2. If you are using an NT account to access the SQL Server database, you must use an NT domain account with appropriate privileges. Additional factors should be considered. See "Database Manager SQL Server Access and Permission Requirements" in the *IBM Tealeaf Databases Guide*.

When NT authentication is enabled, you can configure the Portal to auto-create new Tealeaf user accounts when authenticated NT users first contact the Portal. See "Assigning Tealeaf Product for New NT Users" on page 338.

Figure 93. NT Authentication login screen when Mixed Mode is enabled

Figure 94. Portal Authentication dialog when Mixed Mode is enabled

Figure 95. NT Authentication login screen when Mixed Mode is disabled

Support for Sub-Domains

NT authentication supports the following formats of NT domain user names:

- The full domain:
  `nt.tealeaf.com\thisuser`
- The subdomain only:
  `nt\thisuser`
- The domain only:

```
tealeaf\thisuser
```

Configuring Active Directory

To enable NT authentication use by Tealeaf, some additional configuration within Active Directory may be required. See "Configuring the Search Server" in the *IBM Tealeaf CX Configuration Manual*.

Logging In

After you have enabled NT authentication, you may need to configure additional permissions for your user account or accounts.

- See “CX User Administration” on page 165.

Logging Out

When NT authentication is in use, the Portal can be configured to allow authorized users in the domain to automatically login when they visit the Portal login page.

- In the IBM Tealeaf CX Settings, the `Automatic NT User Login` setting must be set to `Enabled`. See “CX Settings” on page 192.

NT users who have logged in may log out:

- when no action is taken for the period of time, the user is automatically logged out.
- manually by clicking the Logout button.

When NT users automatically log out, they can login manually back into the Portal by visiting the Portal login page. This manual login enables NT users to optionally use a non-NT (Database) account to log into the Portal assuming mixed-mode authentication has been enabled.

- Subsequent logins after a successful login revert to the auto-login method for the user.

Mixed Mode

When NT authentication is enabled, you can optionally enable Mixed Mode authentication, which is useful for performing initial administrative tasks.

- In addition to authenticating the standard NT credentials, Mixed Mode supports authentication using the Portal user credentials.

  Note: Mixed Mode is intended for user management in an NT environment if an NT account with sufficient permissions is not available. It should not be used for other Portal features or general Portal use after accounts have been created or otherwise managed.

Note: Tealeaf Mixed Mode authentication is not related to the mixed mode authentication method available in SQL Server.

Under NT Authentication, users authenticating with Portal credentials have some restrictions; they cannot retrieve events, cannot search, and other restrictions. In Mixed Mode, these same restrictions still apply, yet Portal users with the correct administration permissions can manage Portal users and their permissions. By using Mixed Mode, Tealeaf administrators can grant to NT users the appropriate permissions to use and manage the Portal.

- For example, suppose the Tealeaf administrator has left the company and his NT account is disabled. You can use Mixed Mode authentication to login using the default administrator account (`ADMIN`) and re-assign menu and group permissions to a new administrator.

Note: Mixed Mode authentication is not intended as an additional mechanism for end users to use. It serves an administrative function for managing the initial NT accounts that are later assigned the necessary administrative permissions to manage users.

Web Services Authentication

Like the Portal application, the Portal Web Service (PortalAPI) supports all three methods of authentication: Portal (accounts defined in the Portal), NT (NT/Active Directory accounts) and Single Sign On. Since the Portal and Portal Web Service are both ASP.NET applications, authentication methods for both can be set in similar ways.

However, the Portal Web Service does not have a user interface to enter user credentials, so one of three methods must be used to enable Portal authentication for the Portal Web Service:

- Modify the `AnonAccessPortalAccount` value in the `Web.Config` file to take the Portal account username. This setting enables Anonymous mode for the Portal Web Service, which allows execution of all requests as the specified user.

  Note: This field takes the internal Portal account username, not the User ID.
- Clear the `AnonAccessPortalAcount` value in the `Web.Config` file. If NT/SSO credentials cannot be established and `AnonAccessPortalAccount` is not set, initial connections to the Web Service will be unauthenticated and will not be allowed to go through. At this point, the user must authenticate via the /Authenticate method which accepts a Portal username and password.
  - This mode works if the Portal Web Service is being called programmatically but is an inconvenience if using the Portal Web Service to provide a replay URL to end users.
- Set `ENABLE_PORTAL_AUTH_DIALOG` in the `Web.Config` file to `true`. When an authenticated request is made, a username/password dialog is displayed in the browser, where the authenticating user can enter his or her Portal credentials.

Note: The configuration mode for the Portal Web Service should be set depending on the use case.

Single Sign On (SSO) Authentication

The Tealeaf Portal supports authentication integration with most third-party providers via HTTP request headers. When the SSO mechanism supplies an authenticated username, the Portal performs the login and allows access to the Portal. The Portal relies on the SSO authentication method to ensure the Portal is protected from unauthenticated users and that the SSO headers are valid and unaltered.

- To the Portal, an SSO login is the same as a Portal authentication login.
- SSO users can manually logout and then login again using different credentials. See “Logging Out” on page 332.

Note: Tealeaf officially supports SiteMinder from Computer Associates. Since SiteMinder does not provide a header describing the NT/AD domain from which a user came, the Portal and the search servers cannot be configured to use NT authentication if SSO via SiteMinder is to be used.

The flag to enable SSO support and the header configuration is in the following file `TeaLeaf/Portal/WebApp/Web.config`

```
<configuration>
  ...
  <appSettings>
    ...
    <add key="ENABLE_SITEMINDER_SSO" value="false"/>
    <add key="SITEMINDER_USERNAME_HEADER" value="SM_USER"/>
    <add key="SITEMINDER_SESSIONID_HEADER" value="SM_SERVERSESSIONID"/>
    <add key="SITEMINDER_LOGOUT_URL" value=""/>
    ...
  </appSettings>
  ...</configuration>
```

Note: The configuration options below can be used to provide the required configuration values for other SSO providers.

Property
: **Value**

ENABLE_SITEMINDER_SSO
: The flag to enable (true) or disable (false) SSO/SiteMinder integration.

SITEMINDER_USERNAME_HEADER
: The HTTP request header that specifies the account/user name. The default Tealeaf value is also the SiteMinder default: `SM_USER`.

SITEMINDER_SESSIONID_HEADER
: The HTTP request header that specifies the SiteMinder session ID. The default Tealeaf value is also the SiteMinder default: `SM_SERVERSESSIONID`.

MD5/RC2/3DES/NT
: For more information on supported encryption methods, see "Portal (Database) Authentication" on page 329.

SITEMINDER_LOGOUT_URL
: The HTTP request header that specifies the SSO logout URL. If supplied, the Portal logout button redirects the user to the specified URL to end the current SSO session. If the value is left blank (default), the logout button is hidden, as the Portal cannot end an SSO connection. For the Portal to successfully read the user's NT/AD credentials, IIS must have anonymous access disabled and integrated windows authentication enabled. From the Windows Start menu, these settings can be configured here: **Settings** > **Control Panel** > **Administrative Tools** > **Internet Information Services (IIS) Manager** > **HOSTNAME (local computer)** > **Web Sites** > **Default Web Site** > **Portal** > **Properties** > **Directory Security** > **Edit (Authentication and access control)**. See screen below.

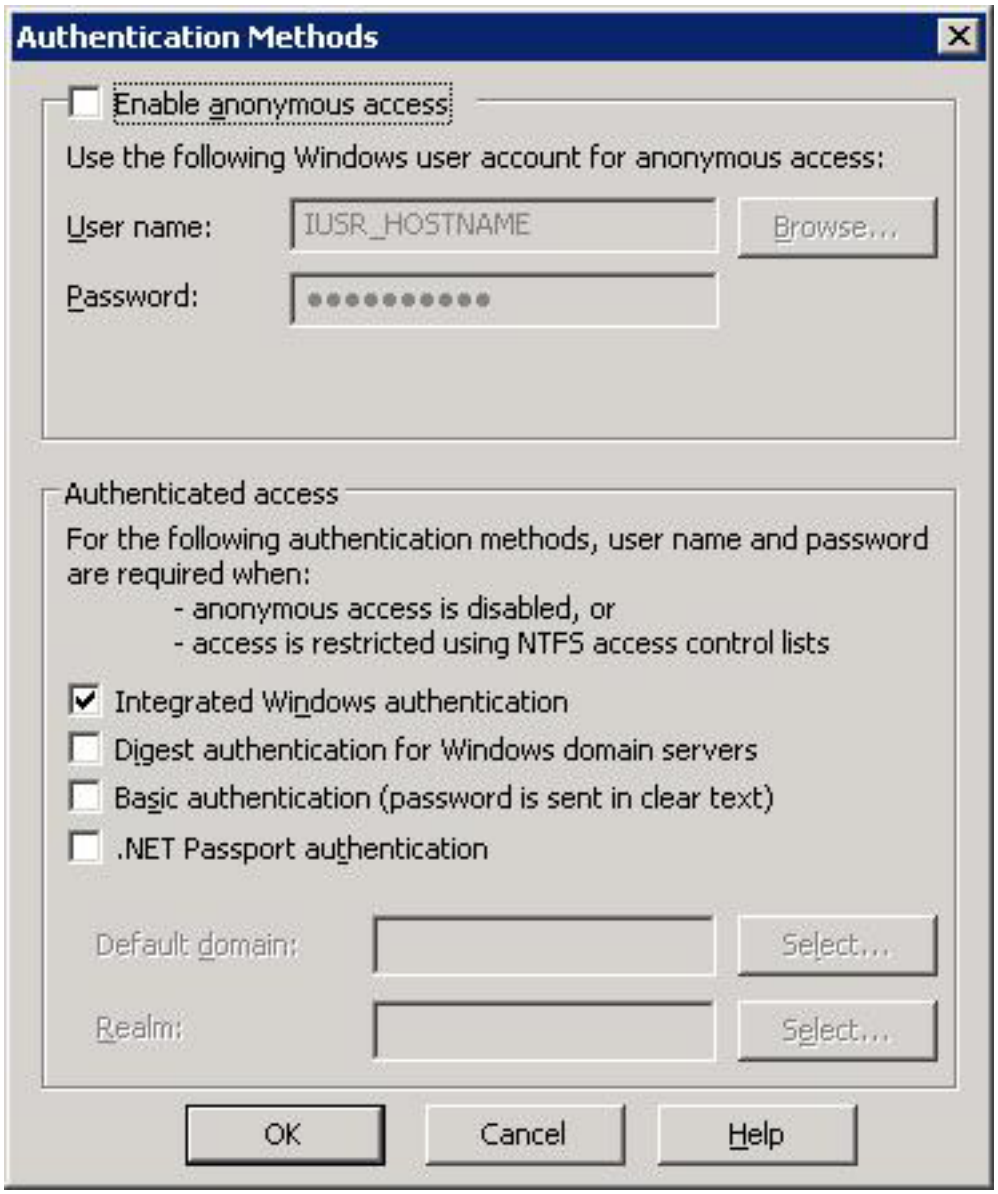

Figure 96. IIS Virtual Directory Authentication Methods

- SiteMinder is a registered trademark of Computer Associates International, Inc.

Additional Required Configuration on IIS 7 for Windows Server 2008

This section describes the additional configuration that is required for IIS 7 for Windows Server 2008.

For NT Authentication

If the Tealeaf Portal application is hosted on IIS 7 for Windows Server 2008, the Windows Authentication scheme may be missing from the list of role services in Server Manager. Please complete the following steps to enable Windows authentication.

Note: This procedure requires an IIS restart, which boots out all Tealeaf users currently logged into the Portal. You should consider performing this procedure during off-peak hours.

1. Login to the server hosting the Portal.
2. From the Start menu, select **All Programs** > **Administrative Tools** > **Server Manager**.
3. In Server Manager, open the Roles node.
4. Right-click the Web Server (IIS) node and select **Add Role Services**.
5. In the Role Services tree, select the Security node.
6. Select the Windows Authentication node.

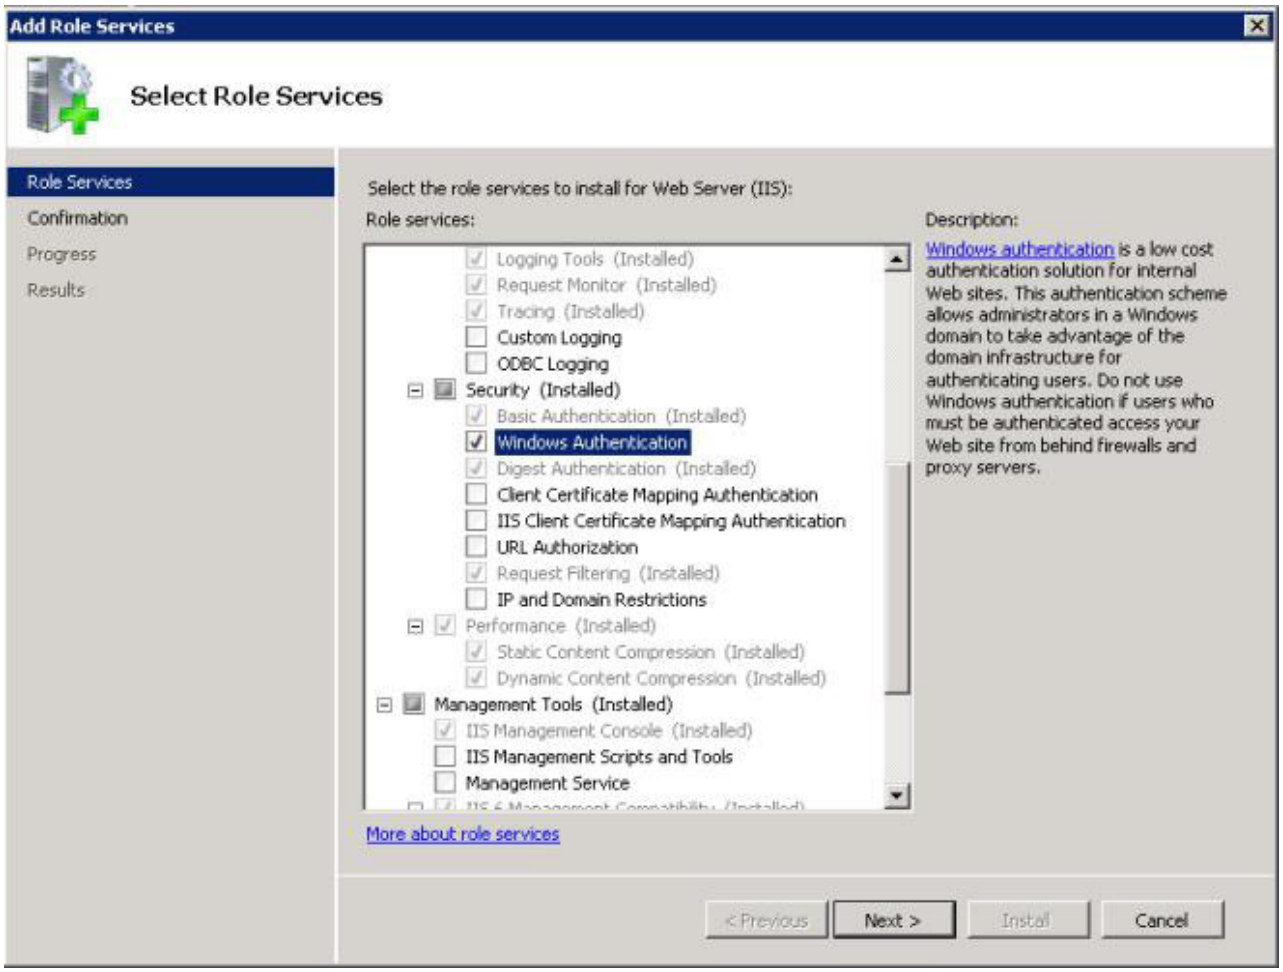
Add Role Services
Select Role Services
Role Services
Confirmation
Progress
Results
Select the role services to install for Web Server (IIS):
Role services:
Logging Tools (Installed)
Request Monitor (Installed)
Tracing (Installed)
Custom Logging
ODBC Logging
Security (Installed)
Basic Authentication (Installed)
Windows Authentication
Digest Authentication (Installed)
Client Certificate Mapping Authentication
IIS Client Certificate Mapping Authentication
URL Authorization
Request Filtering (Installed)
IP and Domain Restrictions
Performance (Installed)
Static Content Compression (Installed)
Dynamic Content Compression (Installed)
Management Tools (Installed)
IIS Management Console (Installed)
IIS Management Scripts and Tools
Management Service
Description:
Windows authentication is a low cost authentication solution for internal Web sites. This authentication scheme allows administrators in a Windows domain to take advantage of the domain infrastructure for authenticating users. Do not use Windows authentication if users who must be authenticated access your Web site from behind firewalls and proxy servers.
More about role services
< Previous
Next >
Install
Cancel

Figure 97. Adding Windows Authentication role service in IIS 7

7. Click **Next**.
8. Review the Confirmation information. To install, click **Install**.
9. After the service is installed, click **Close**.
10. A restart is not required.
11. Open Internet Information Services (IIS) Manager.
12. Select the server node. Click **Sites** to display the sites hosted by the server. The Tealeaf Portal should be one of them.
13. In the Tealeaf web site node, click **Portal**.
14. Double-click the Authentication icon.

15. Right-click **Windows Authentication** and select **Enable**.

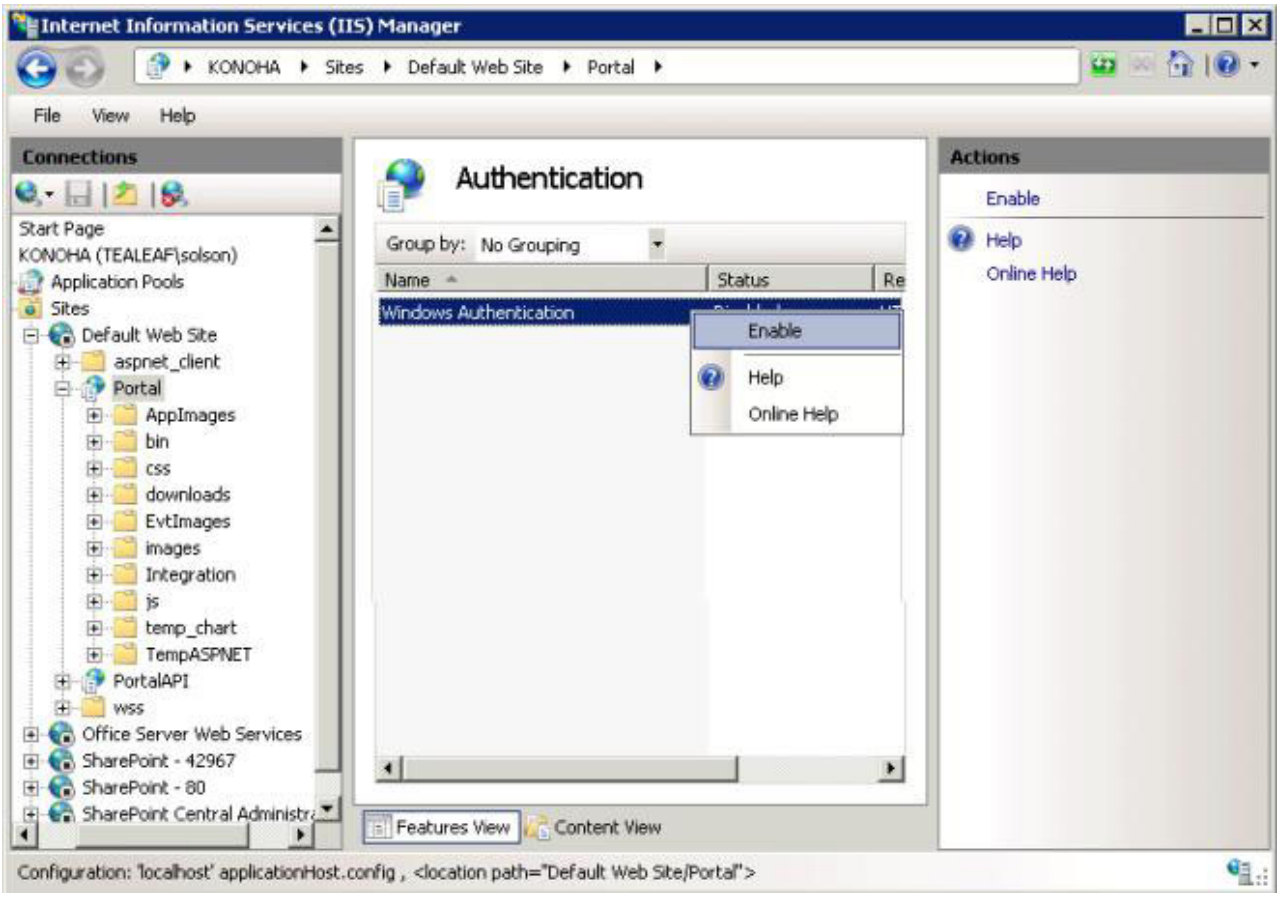

Figure 98. Enabling Windows Authentication in IIS 7

16. Windows authentication is enabled.
17. If Forms Authentication is disabled for the Portal, right-click and select **Enable**.
18. If Anonymous Authentication is enabled for the Portal, right-click and select **Disable**.
19. Restart the Tealeaf Data Service.
20. When the Data Service has restarted, perform an IIS reset.
21. If NT Authentication has already been enabled through the Portal, verify that the Portal is accessible using NT credentials.
    - See “NT (AD - Active Directory) Authentication” on page 330.

Note: If you have enabled the Portal Web Service, in the most common case, you must enable access for Windows Authentication and Anonymous Authentication.

For non-NT Authentication

If you have not enabled NT authentication for IIS, please verify that the following settings are configured.

1. Open Internet Information Services (IIS) Manager.
2. Select the server node. Click **Sites** to display the sites hosted by the server. The Tealeaf Portal should be one of them.

3. In the Tealeaf web site node, click **Portal**.
4. Double-click the Authentication icon.
5. Right-click **Windows Authentication** and select **Disable**.
6. If Forms Authentication is disabled for the Portal, right-click and select **Enable**.
7. If Anonymous Authentication is disabled for the Portal, right-click and select **Enable**.
8. Restart the Tealeaf Data Service.
9. When the Data Service has restarted, perform an IIS reset.
10. Verify that the Portal is accessible using Portal credentials.

Configuring Automatic User Creation in NT Authentication

If NT authentication is enabled for Portal access, you can configure the Portal to automatically create a user account when an authenticated user first contacts the Portal. Depending on the Tealeaf products you have licensed, you can determine the type of user to create, the account's default permissions, and any group assignments.

Enabling NT Authentication

If you have not done so already, you can enable NT authentication for the Portal. When enabled, the Portal relies on the NT domain user permissions to enable access to the Tealeaf Portal. See "NT (AD - Active Directory) Authentication" on page 330.

Assigning Tealeaf Product for New NT Users

Through the Portal Management page, you can enable automatic user creation globally and specify the licensed Tealeaf product to which first-time users who are NT-authenticated are assigned.

1. Login to the Tealeaf Portal as an administrator.
2. In the Portal menu, select **Tealeaf** > **Portal Management**.
3. In the left navigation bar, click the IBM Tealeaf CX Settings link.
4. Click the Users link.
5. In the IBM Tealeaf cxImpact Users settings, the following settings must be specified in order to enable automatic user creation.

Table 35. Assigning Tealeaf Product for New NT Users

Setting	Description	Assign This Value
`Automatic User Creation`	When set to `Enabled`, first-time visitors to the Tealeaf Portal are automatically given a user account, whose username is their NT domain login ID.	`Enabled`
`Automatic User Creation Settings Prompt`	When set to `Enabled`, automatically created users are allowed to change their settings upon login.	`Disabled`
`Automatic User Creation Type`	This setting defines the Tealeaf product to which automatically created users are initially assigned: `Portal` (base IBM Tealeaf cxImpact account), `cxReveal`, or `cxView` account	Depends on your Tealeaf solution

Table 35. Assigning Tealeaf Product for New NT Users (continued)

Setting	Description	Assign This Value
`New User Default Group`	The default user group to which new users are automatically assigned. This value should correspond to the value assigned for `Automatic User Creation Type`. For example, if new users are `cxReveal` type, the default new user group should be a IBM Tealeaf cxReveal group.	Depends on your Tealeaf solution
`New User Default Page`	The page that is displayed to new users after login. **Note:** The new user must have permissions to access the selected page. You should verify that the page is accessible. For more information, see the Group settings on the User Administration page for the product: • "CX User Administration" on page 165 • "cxReveal User Administration" in the *IBM Tealeaf cxReveal Administration Manual* • "cxView User Administration" on page 184	Depends on your Tealeaf solution

6. For security reasons, you may wish to define the password policy in advance of generating automatic user accounts. Please review the following settings:

   Setting Description

   `Previous Password History (Count)`
   When users are prompted to set a new password, they may be prevented from using a password that has been set within the number of times defined in this setting. To disable this check, set the value to `Disabled`.

   `Previous Password History (Days)`
   When users are prompted to set a new password, they may be prevented from using a password that has been set within the number of days defined in this setting. To disable this check, set the value to `-1`.

   `Require Strong Passwords`
   When `Enabled`, users are required to assign a password that passes Tealeaf's strong password test. To disable this check, set the value to `Disabled`.

For more information on these settings, see "CX Settings" on page 192.

Configure User Auto-Creation for Individual Products

After automatic user creation has been enabled, you must configure the auto-create user settings for the individual product to which users are first assigned.

- See "CX Settings" on page 192.
- See "cxReveal Settings" in the *IBM Tealeaf cxReveal Administration Manual*.

Custom HTML content for the login page

File

This file is dynamically loaded into the section beneath the Portal login box on the login page. In the first line of the table definition below, the default HTML has a

'display: none' style applied.

```
<table  class="bordered" border="0" cellpadding="3" cellspacing="0">
  <!-- Remove 'display: none' style property before use -->
  <tr class="customContentHeader">
    <td align="center">
      <!-- Place Custom Content Here -->
      Place content here.
      <!-- End Custom Content -->
    </td>
  </tr></table>
```

Figure 99. custom_content.htm

To enable custom HTML content, you need to remove the `display : none` style.

To increase performance, the Portal caches this file in memory. To refresh the cached version of this file, add a `resetcustomcontent=1` URL parameter to the URL title bar. Examples:

Figure 100. Custom content on the login page

Custom Menu Group

File

`TeaLeaf\Portal\WebApp\custom_menu.xml`

This file can be modified to add to the Tealeaf Portal your customized menu group with one level of child items. The `custom_menu.xml` supports one CustomMenu element, which supports one Title element and the Items list element. The Items list element supports both Item and Divider elements.

CustomMenu Element

The CustomMenu element's `enabled` attribute enables or disables the entire custom menu group. The `firstgroup=true/false` CustomMenu attribute controls whether the custom menu group displays as the first group on the menu (far left) or the last group on the menu (far right).

Items Element

The Items list element is simply a container for the child menu items.

Item Element

The Item element supports the following child elements.

Element
Description

Title (Required) The text to display on the menu item

URL The relative or absolute URL to navigate to JavaScript

JavaScript
JavaScript snippet to execute on the client

Either URL or JavaScript should be provided. If both are present, URL value is used.

Item elements can independently be enabled or disabled from one another.

Divider Element

The divider element renders a horizontal divider to separate groups of items.

Just like the Item element, Divider elements can be enabled or disabled with no impact to any other Divider of Item elements.

```
<CustomMenu enabled='true' firstgroup='false'>
  <Title>Custom Menu Group</Title>
  <Items>
    <Item enabled='true'>
      <Title>Menu Item 1</Title>
      <URL>./About.aspx</URL>
    </Item>
    <Divider enabled='true'/>
    <Item enabled='true'>
      <Title>Menu Item 2</Title>
      <URL>./About.aspx</URL>
    </Item>
    <Item enabled='true'>
      <Title>Menu Item 3</Title>
      <URL>./About.aspx</URL>
    </Item>
  </Items></CustomMenu>
```

Figure 101. custom_menu.xml

To increase performance, the Portal caches this file in memory. To refresh the cached version of this file, add a resetcustomcontent=1 URL parameter to the URL title bar while developing. Example:

```
http://machine/Portal/About.aspx?resetcustommenu=1
```

The cache auto-expires in five minutes. If there is no need to immediately refresh the custom_menu.xml file, replace the existing file. Within five minutes, the Portal automatically loads the new file.

Figure 102. Custom menu group

Character Indexing

Searching for sessions in the Long-Term Canister utilizes an embedded search engine called dtSearch, which relies on an index of words in captured sessions. When a search is queried, dtSearch scans these indexes of to quickly find sessions that match search criteria.

Overview

For the dtSearch engine, the base unit of indexing is a word. For example, to search for the word `apple`, you must enter the full word `apple` in Portal search.

- To search for part of the word, you must use a wildcard character (`*`). To find `ap`, enter `ap*`.

To index words, dtSearch must understand which characters constitute words and which constitute punctuation such as whitespace, and which constitute noise. dtSearch only indexes words composed of characters that it doesn't consider to be noise.

Alphabet.dat

dtSearch defines what constitutes a word using an alphabet file, which designates the characteristics of each character that can be encountered in session data. This file is located here:

```
<Tealeaf_install_directory>\Alphabet.dat
```

`Alphabet.dat` file is a text file. For characters that are not displayed in text files, ASCII equivalents may be expressed using the slash and the hexadecimal code for the character. For example, `\0c` is the code for ASCII character 13, which is a return key.

- When an index is created, the dtSearch engine stores a private copy of the current alphabet and hyphenation settings in the `index_a.ix` file in the index folder. These settings are used to index all files added to the index and to evaluate search requests for the index.

Character Categories

dtSearch classifies characters into four categories: letter, space, hyphen, and ignore. These categories are specified in the headings listed below:

Table 36. Character Categories

Heading	Category	Meaning
[Letters]	letter	A searchable character. All of the characters in the alphabet (a-z and A-Z) and all of the digits (0-9) should be classified as letters. The alphabet specifies whether a letter character is lowercase or uppercase, whether or not it has an accent, and the lowercase or unaccented equivalent.
[Hyphens]	hyphen	Hyphen characters can receive special processing in dtSearch. By default, only the - character is defined as a hyphen. • You can configure how hyphens are interpreted for indexing purposes.
[Spaces]	space	A character that causes a word break. For example, if you classify the period as a space character, then dtSearch would process U.S.A. as three separate words: U, S and A.
[Ignore]	ignore	A character that is disregarded in processing text. For example, if you classify the period as ignore instead of space then dtSearch would process U.S.A. as one word: USA.

The file must be concluded with an ending heading:

```
[End]
```

Replaced Characters

During indexing, the following non-alphanumeric characters are replaced with the underscore character (_).

```
[ ] . $ :
```

When indexing field values or text outside fields, non-alphanumeric characters are treated according to their category in `alphabet.dat` file. When indexing field names, the above characters are replaced with the underscore.

For example, if the request field and value was the following:

```
[urlfield]
paymentDetailsVO.title=foo.
ctl01$PlaceHolderMain$ctl00$=bar.
```

The text indexed for search is the following:

```
<urlfield>
  <paymentDetailsVO_title>foo.</paymentDetailsVO_title>
  <ctl01_PlaceHolderMain_ctl00_>bar.</ctl01_PlaceHolderMain_ctl00_>
</urlfield>
```

Treatment of special characters before submitting a search query

The following characters are stripped from Portal-based search criteria and replaced with spaces before submitting to dtSearch.

\ / : & < > % ! ? . @ # $ ^ ( ) ? { } | ' ~ , * [ ] =

The asterisk character (*) is not stripped from the field, since it is used as a wildcard character.

Scheduling Tealeaf Jobs

The Tealeaf Scheduling Service is a Windows service through which you can create and manage Tealeaf jobs. You can configure jobs, which the service then runs from the command line at the scheduled intervals.

Note: Do not use the Tealeaf Scheduling Service for non-Tealeaf jobs.

To configure Tealeaf jobs:

1. Verify that the Tealeaf Scheduling Service is automatically started through the Services Control Panel.
2. You can configure and enable individual jobs through the following XML file:
   `<Tealeaf_Install_Directory>\Tools\TLSchedulerSvcCfg.xml`
3. The following types of jobs can be scheduled through the service:
   a. **`TLBackup.exe`** Cycle the Tealeaf services. This job should be scheduled to run once per day during off-peak hours.
   b. **`TLExtractor.exe`** This job runs the data extractor on the Tealeaf Visitor database, if it is installed in your Tealeaf solution. It can be executed once every fifteen minutes.

See "Configuring the Scheduling Service" in the *IBM Tealeaf CX Configuration Manual*.

Scheduling Overnight Tealeaf Tasks

During off-peak hours, some maintenance tasks must be performed or should be performed to the Tealeaf system, its Canister data, and the databases. This section provides a list of the most important Tealeaf tasks and a suggested schedule for their execution.

Note: This list of tasks does not include any tasks that should be performed by your enterprise systems, such as database backup and virus scanning. This schedule should be reviewed with your IT staff.

Note: If you have any questions about this schedule, please contact Tealeaf Professional Services.

Suggested Scheduling of Overnight Tasks

Note: For purposes of the following schedule, it is assumed that the workday begins at 6am and ends at 6pm. All scheduling tasks are configured on based on the Tealeaf system time. See "Configuring the System Timezone" in the *IBM Tealeaf CX Configuration Manual*.

In the table below, time is listed on the 24-hour clock.

Table 37. Suggested Scheduling of Overnight Tasks

Task	Product	Component	Description and Dependencies	Time
Cycle Services	IBM Tealeaf cxImpact	Scheduling Service	All Tealeaf services should be cycled once. See "Configuring the Scheduling Service" in the *IBM Tealeaf CX Configuration Manual*.	0:30
TLBackup	IBM Tealeaf cxImpact	Scheduling Service	Tealeaf provides the TLBackup utility to backup Canister data. See "Configuring the Scheduling Service" in the *IBM Tealeaf CX Configuration Manual*.	0:30
Daily Data Aggregation	IBM Tealeaf cxImpact	Data Collector	Once per day, daily reporting data is aggregated for insertion into the Reporting database. • This task must be completed before the Data Collector's trimming task. • See "CX Settings" on page 192.	2:00
TLI Trim	IBM Tealeaf cxImpact	Scheduling Service	If a TLI server has been deployed, the TLI trim task removes static content from the TLI server after it has aged a preconfigured number of days. See "Configuring the Scheduling Service" in the *IBM Tealeaf CX Configuration Manual*. **Note:** TLI Trim and TLI Merge should be scheduled for the same time. • For more information on TLI, see "Managing Static Archives" on page 374.	2:30
TLI Merge	IBM Tealeaf cxImpact	Scheduling Service	If a TLI server has been deployed, the TLI merge task merges the daily TLI files into the monthly TLI files for longer-term storage. See "Configuring the Scheduling Service" in the *IBM Tealeaf CX Configuration Manual*. **Note:** TLI Trim and TLI Merge should be scheduled for the same time. • For more information on TLI, see "Managing Static Archives" on page 374.	2:30
Data Trimming	IBM Tealeaf cxImpact	Data Collector	Periodically, the Data Collector trims aged data from the reporting database. This option is applied if data trimming has been set to `Daily` or `Monthly`. • See "CX Settings" on page 192.	3:00
Dimension Trimming	IBM Tealeaf cxImpact	Data Collector	On a nightly basis, the Data Collector trims dimension values if the number of values exceeds the preconfigured maximum number of them and if dimension trimming is set to `Daily`. • The default is `Hourly`. See "CX Settings" on page 192.	3:00
Dimension Log Aggregation	IBM Tealeaf cxImpact	Data Collector	When dimension log aggregation is configured to occur on a daily basis, the Data Collector reviews dimension logs for aggregation of reporting data. See "CX Settings" on page 192.	3:00
Top Movers	IBM Tealeaf cxImpact	Data Collector	When Top Movers are enabled, the Data Collector performs calculations of Top Movers at the defined time. See "CX Settings" on page 192.	4:30

Table 37. Suggested Scheduling of Overnight Tasks (continued)

Task	Product	Component	Description and Dependencies	Time
Database Growth	IBM Tealeaf cxImpact	Data Collector	Calculates the growth of each database for purposes of Portal reporting. See “System Status” on page 8.	5:00

Other Tasks

Table 38. Other Tasks

Task	Product	Component	Description and Dependencies	Default Schedule
Tealeaf Status	IBM Tealeaf cxImpact	Scheduling Service	Tealeaf Status gathers information from Tealeaf servers known to the Portal about their current state. See "Configuring the Scheduling Service" in the *IBM Tealeaf CX Configuration Manual*.	Every 15 minutes
Visitor Database Extractor	IBM Tealeaf cxResults	Scheduling Service	The Visitor Database Extractor task queries the Canisters for session and event data to insert into the Visitor Staging database for use by IBM Tealeaf cxResults. See "Configuring the Scheduling Service" in the *IBM Tealeaf CX Configuration Manual*.	Every 15 minutes
Extract Service	IBM Tealeaf cxConnect for Data Analysis/IBM Tealeaf cxVerify	Extract Service	If you have licensed and deployed IBM Tealeaf cxConnect for Data Analysis or IBM Tealeaf cxVerify, the Extract Service is used to extract session data on-demand for delivery to third-party systems. This service must be enabled. See "Configuring the Extract Service" in the *IBM Tealeaf CX Configuration Manual*. • For more information on configuring IBM Tealeaf cxConnect for Data Analysis tasks to use this service, see "cxConnect Configuring Tasks" in the *IBM Tealeaf cxConnect for Data Analysis Administration Manual*. • For more information on configuring IBM Tealeaf cxVerify tasks to use this service, see "cxVerify Configuring Tasks" in the *IBM Tealeaf cxVerify Administration Manual*.	On-Demand

TLBackup and TLRestore

Implementing a periodic backup schedule protects you from data loss and keeps a record of important sessions for future reference. Although backing up and removing session data on a daily basis is the ideal scenario, you should implement a backup plan that is appropriate for your production environment.

Note: These utilities perform backups of Canister data only. For more information on backing up other Tealeaf data, see “Backing Up Tealeaf” on page 352.

Depending on the volumes of data you are backing up, you may find it better to perform a partial backup instead of a full backup. A partial backup saves all canister data files that have not been previously archived.

Note: Except for the Tealeaf Transport Service, all Tealeaf Services are stopped during a TLBackup or TLRestore operations. This interruption may result in significant spooling of sessions and should be scheduled during off-peak hours.

Other parameters to consider:

- **Frequency** - The frequency at which you back up data depends on several factors: the number of interesting sessions you have chosen to archive, the amount of traffic your Web site has, and the number of sessions you can afford to lose. If you have a low tolerance for data loss, you should schedule backups every day. You can schedule Canister backups by creating a scheduled task. See "Scheduling Canister Backups" on page 350.
- **Retention** - Determine the length of time you want to keep backup data.
- **Location** - In most cases, you should have a designated backup computer on which to store backups. You can install TLBackup on all Canister machines and specify to back up to the same remote backup machine.

TLBackup can perform both full and incremental backups of canister data and various configuration files.

For more information on configuring TLBackup jobs, see *IBM Tealeaf CX Configuration Manual*.

Coordinating CanTrim and TLBackup in All-in-One Environments

CanTrim deletes sessions from the Canister older than a specified number of days, while backup archives sessions at specified intervals.

- Specify a time for CanTrim to run that is later than TLBackup. If CanTrim runs before TLBackup, it could permanently delete sessions that you intend to archive.
- When TLBackup runs, it clears the archive attribute bit after it backs up LSSN data file. CanTrim checks to see if this bit is cleared. If it has not been cleared, then it does not run. This creates a potential situation where sessions will continue to grow in the Canister and consume more disk space if backup is not run regularly. If you are not running backup regularly and want to verify that sessions get deleted from the Short Term Canister, deselect the `Delete session data only after backup` checkbox in the Canister Configuration dialog box so CanTrim deletes sessions older than the number of days specified.
- **Full backup** - Copies the entire base Canister to a specified TLBackup directory.
- **Incremental backups** - Only copies the `LSSN*.dat` and `LSSN*.IDX` day data files to a specified TLBackup directory.

Running TLRestore

TLRestore enables you to restore backup directories in the order of restoration, which allows you to rebuild your dataset through a series of Partial backups. TLRestore first asks to perform a rebuild of the Canister, which removes all current data. When the canister is rebuilt, all backup directories are copied in the appropriate order to the `Canister.dbs` directory.

- When the restore is complete, you must restart the Tealeaf Services and perform a check-and-fix operation to rebuild the indexes. These tasks are described below.

To restore Canister data, perform the following steps:

Note: This process replaces all data from the canister.

Note: This process recreates the indexes, which can take a considerable amount of time. Please schedule this restore during off-peak hours.

1. Run `<tealeaf_install_directory>\tools\TLRestore.exe`.
   - Rebuilds the canister, which deletes all current data in the canister.
   - Restores all data listed in the backup directories in the order of restoration.
2. Run a check-and-fix on the indexes to recreate them.
   a. Enter the following URL on the Search Server:
      `http://<search_server_name>:<port (19000)>/CIC`
   b. Click **Check and Fix**.
   c. This process can take a significant amount of time.
3. Start all Tealeaf services.

Configuring TLBackup

Configuration is accomplished in two basic steps:

- Configuring the parameters in `TLBackupCfg.xml`. These parameters are described below.
- Scheduling when to run TLBackup.

1. To Configure TLBackup, open `TLBackupCfg.xml` file located in the `Tools` directory in the Tealeaf install directory.
2. Modify the following configuration options:

   Option Description

   `Backup Directory`
   Specifies the name of the directory in which to store the Canister data files. This directory can be a remote host on the network. To specify the remote host, use the UNC path naming: `\\BACKUPSERVER\directory`.

   Note: The specified backup directory on the remote server must be shared. The user under which the backup process is scheduled to run must have permissions to write to the designated backup server.

   `Remove Old TLBackup Directories`
   When set to 1, this option removes old backup folders. The `FullBackUpCount` setting specifies the number of backup folders to keep. This option works by timestamp, so only the most recent backup folders are kept.

   `FullBackUpCount`
   Specifies the number of old backup folders to retain.

   `BackupDaysToSave`
   Specifies the number of partial backup folders to retain.

   `FreeDiskSpace`
   Specifies additional free space to allot on the hard drive beyond what is required by `TLbackup.vbs`.
3. To run a single instance of TLBackup, open a command prompt.
4. Navigate to the `Tools` directory inside the Tealeaf install directory.
5. Run `TLBackup.exe`:
   a. You can run this manually at the command line.

b. You can also set up a Windows task to run the script automatically at a specified interval. To run the script, execute the following command:

```
Tlbackup.exe -Partial
```

Note: For this process, you should perform a partial backup, which does perform a full backup of all canister data files that have not been previously archived.

6. The following options are available through the command line or through the Windows Task Scheduler:

Option Description

Full Specifies to back up the base Canister. This includes the Master Session archive table (LSSN), Post Event/ Annotation table (PEVT), Report tables (RF*.DAT), Event and Category tables (NTYP, NCTG), corresponding CSV files, and Search Server registry files. See “Full Backup” on page 351.

Partial
Specifies to back up the incremental Canisters with the archive field value set. This backs up LSSN*.DAT and corresponding .IDX files, PEVT*.DAT and corresponding .IDX files. See “Partial Backup” on page 351.

Cycleservices
Specifies to stop and restart Tealeaf services. Before stopping services, TLBackup verifies that sessions are drained from the Short Term Canister before proceeding with services shutdown. Upon services restart, TLBackup starts TLTMaint to perform an integrity check on the tables. Progress is logged to the console window. See “Cycle All Services” on page 351.

TLBackup performs the following checks on startup

- Retrieves time zone and calculates a day based on the time zone configured at install time.
  - See "Configuring the System Timezone" in the *IBM Tealeaf CX Configuration Manual*.
- Derives Hostname of indexing computer from the Windows Registry if the IBM Tealeaf CX Indexer is not on the same machine.
- Reads the ctsrvr.cfg files to get the location of the Tealeaf datastore.
- Checks for available disk space.
- Stops services and run TLTMaint to check the integrity of the Canister. TLTMaint launches CanTrim, if it finds sessions older than the specified number of days.
- Performs backup operation. No data is lost during backup as incoming data is queued by the Extended Decoupler.
- Resets the archive bit.
- Removes old backup folders if specified.

Format of the TLBackup directories

When TLBackup.exe is run, it automatically creates and stores backed up data into subfolders named with the year, date, time, and a prefix of _FULL or _PARTIAL. Example: 200207291648_FULL.

Scheduling Canister Backups

Use the Windows Scheduled task wizard to automate all backup tasks.

Note: If a backup is required, Tealeaf recommends running a Partial backup.

Scheduling Daily Full Backups of the Long Term Archive

1. Open Windows Explorer. Click the My Computer node in the left pane to expand the folders tree.
2. Click the Control Panel folder.
3. Choose the Scheduled Tasks folder from the list.
4. Double-click **Add Scheduled Task** to start the Scheduled Task wizard.
5. Click **Next** and then click **Browse**. Select the following:
   `<Tealeaf_install_directory>\Portal\Tools\TLbackup`
6. On the first page of the wizard, enter a name for the task. It is helpful to add the type of backup. For example: `Partial TLBackup`.
7. Choose the frequency with which you want to perform the backup. For a Daily backup, choose the `daily` option.
8. Select a start time. Choose a time when Web traffic is at its lowest. This is typically early morning hours. For example, 3:00 AM.
9. In the Task tab in the Run field, enter the `-Partial` parameter after the path and TLBackup file name.

   Note: TLBackup must have a command-line parameter to initiate the backup process. For a list of command line parameters, “Configuring TLBackup” on page 348.
10. When prompted for the user and password to execute the task, enter a username that has full write permissions on the Canister and backup machines.

Restoring Tealeaf Data from TLBackup

`TLRestore` copies data files for the base canister and incremental day to the current canister directory.

Note: When you perform a restoration of a Canister backup, all data stored in the Canister is lost and cannot be recovered.

RestoreCanister reads the `TLBackupCfg.xml` file to locate the backup directory and to determine the settings used for the original backup.

To Restore theIBM TealeafCX Datastore:

1. Open a command prompt and navigate to the following directory:
   `<TeaLeaf_install_directory>\Portal\Tools\TLbackup`
2. Open `TLRestore.exe`.
3. TLRestore runs through the following list of processes:
   - RestoreCanister locates the Tealeaf directory
   - Rebuild Canister
   - RestoreCanister restores the base Canister
   - Full restore or partial restore

Command-Line Options

In addition to backing up Tealeaf, TLBackup is the service used to cycle services in the Tealeaf system, except for the Transport Service, which must remain present to spool data during the TLBackup operation.

- These options can also be configured through the XML configuration file for TLBackup. See "Configuring TLBackup" on page 348.

You can execute backup and cycle services using the following commands at the command line.

Partial Backup

```
tlbackup -Partial
```

Full Backup

```
tlbackup -Full
```

Cycle All Services

```
tlbackup -Cycleservices
```

If a Tealeaf service is in a stopped state when TLBackup is executed, TLBackup does not start the service. If you have intentionally taken a service down for maintenance or troubleshooting, then this behavior prevents an unwanted restart of the service in maintenance.

To force a restart of all services after TLBackup has executed, you can issue the following command-line command. This command should be executed after TLBackup.exe has completed, preferably as a second command in the same command-line script:

```
"<tealeaf_install_directory>\CanSvcs.exe" -start
```

where:

- `<tealeaf_install_directory>` is the full path to the Tealeaf install directory.
- The double-quotes are required.

When the above command is executed, any running services are not affected, and any stopped services are started.

Use Alternate Configuration File

You may specify an alternate configuration file in the command, which enables you to perform different types of backups depending on the situation, without modifying a single config file between backup operations.

In the example command below, a full backup is performed using the alternate config file `[cfgfile]`:

```
tlbackup -Full -Altcfg [cfgfile]
```

here:

`[cfgfile]` is the full path to the alternate configuration file.

Get Command Line Help

```
tlbackup -?
```

Backing Up Tealeaf

Tealeaf provides multiple mechanisms for backing up your Tealeaf system and your captured data.

Backup Type
Description

Event Model data
Through the Portal Management page, you can execute scheduled and on-demand backups of event model data. For more information, see "Backing up the IBM Tealeaf event model" on page 254.

Database data
Using the Tealeaf Database Manager or SQL Server Management Studio, you can perform full backups of all Tealeaf databases.

- For more information on database backups in general, see "Database Backup Strategy" in the *IBM Tealeaf Databases Guide*.
- For more information on performing backups through the Tealeaf Database Manager, see "Using Tealeaf Database Manager" in the *IBM Tealeaf Databases Guide*.

Canister data
Tealeaf provides a pair of utilities for backing up session and index data stored on Tealeaf Canisters. For more information, see TLBackup and TLRestore.

scorecards/dashboards/reports
Through the Portal, you can import and export templates for scorecards, dashboards, and reports. For more information, see "Importing and Exporting Reports" in the *IBM Tealeaf cxView User Manual*.

Configuring Maximum Capture Sizes for Pages, Hits, and Sessions

This section describes how you can change the maximum request size and maximum response size that Tealeaf is permitted to capture.

By default, Tealeaf is configured to capture requests and responses up to a predefined size. This default limit is intended to prevent the capture of large hits.

In most implementations, the default maximum capture size is sufficient for handling all useful web application traffic. If needed, you can change the maximum permitted size for individually captured hits.

Limits on hit data size are defined in multiple locations in the Tealeaf software.

Monitoring Oversized Hits

This section describes how PCA and pipeline hits are logged when they exceed the maximum permitted size.

Monitoring for PCA

This section describes how PCA hits are logged when they exceed the maximum permitted size.

Response size

When the response portion of a hit exceeds the maximum permitted size, the response is truncated to the max size, and a log entry similar to the following is inserted:

```
TLAPI: Truncating body because RequestLength(%u) > MaximumRequestLength(%u)
for connection: %s
```

For more information on changing the maximum size, see "Whole hit size."

Request size

When the request portion of a hit exceeds the maximum permitted size, the hit is dropped, and a log entry similar to the following is inserted:

```
TLAPI: Dropping hit because RequestHeaderBytesRead(%u) >=
MaximumRequestLength(%u) for connection: %s
```

The default value for the Max Request Size setting is 1.5MB, or 1572864 bytes.

Whole hit size

When a request and its associated response (called **a hit**) exceeds the maximum permitted size, it is dropped.

Before the hit is sent from the PCA to its delivery target the size of the hit is checked. If the size exceeds the maximum data size allowed (2MB by default), the hit is dropped, and a log entry is similar to the following is inserted:

```
TLAPI: Data size (%u) exceeds default MaxDataSize (%u) for connection: %s
```

Monitoring for Tealeaf Pipeline

When a pipeline hit exceeds the configured maximum size, an Event Log entry is generated, and its message indicates the size of the offending hit.

Configuring Max Data Size

This section describes how to configure the maximum data size of PCA and pipeline hits.

Configuring for PCA

In the PCA Web Console, when a hit exceeds the `Max Data Size` value defined (2MB by default), the `Total dropped because hit exceeded max data size` statistic is incremented. This stat is available in the Statistics tab of the Console.

- See [Passive Capture Configuration via Web Console].

HTTP request and response sizes

When the PCA assembles TCP connection data into requests and responses, the following parameters define whether individual requests and responses are monitored for maximum size and whether the connection is closed when the defined limit is exceeded.

In the `<Capture>` section of the `ctc-conf.xml` file, the following settings define and enable the checking of maximum individual request or response size at the TCP connection level. It is the first size check of this type within the PCA software chain:

```
<DeleteTcpLargeConnDisabled>false</DeleteTcpLargeConnDisabled>
    <MaxTcpConnSize>2097152</MaxTcpConnSize>
```

1. If <DeleteTcpLargeConnDisabled> is set to true, no check on the data size is performed. If so, the <MaxTcpConnSize> is not used.
   - The default value is false.

   Note: This value is not exposed in the Web Console and must be configured by manually editing the ctc-conf.xml file. See "Passive Capture Configuration File ctc-conf.xml" in the *IBM Tealeaf Passive Capture Application Manual*.
2. If <DeleteTcpLargeConnDisabled> is set to false or is not present in the config file, then <MaxTcpConnSize> defines the maximum size of a single request or response.
   - The default value is 2MB (2097152 bytes).
   - If the maximum size is exceeded, then the TCP connection is automatically closed.

   Note: This value is not exposed in the Web Console and must be configured by manually editing the ctc-conf.xml file. See "Passive Capture Configuration File ctc-conf.xml" in the *IBM Tealeaf Passive Capture Application Manual*.

Max size during parsing

When the PCA pipeline is parsing requests and responses, three parameters are applied to determine the maximum permitted size of each request, response, and the entire hit.

In the <Parse> section of the ctc-conf.xml file, the following settings are used to define the maximum size during this stage of PCA processing:

```
<MaxRequestSizeBytes>2097152</MaxRequestSizeBytes>
    <MaxResponseSize>1572864</MaxResponseSize>
    <MaxDataSizeBytes>2097152</MaxDataSizeBytes>
```

Setting Description

< MaxRequestSizeBytes>
: If the request size is greater than the maximum value, the hit is dropped. The default is 2 MB.

< MaxResponseSizeBytes>
: The <MaxResponseSize> value is used to check the size of both the request and the response. The default is 1.5 MB.
  - If the response size is greater than the maximum value, the hit is truncated to the maximum value. See “Response size” on page 353.
  - This value is exposed and configurable in the Pipeline tab of the Web Console. See "Passive Capture Configuration via Web Console" in the *IBM Tealeaf Passive Capture Application Manual*.

< MaxDataSizeBytes>
: The <MaxDataSizeBytes> value is used to check the total size of a hit (request and response) before it is sent from the PCA to the target Windows machine. The default is 2 MB.
  - If the hit size is greater than the maximum value, the hit is dropped.

  Note: This value is not exposed in the Web Console and must be configured by manually editing the ctc-conf.xml file. See "Passive Capture Configuration File ctc-conf.xml" in the *IBM Tealeaf Passive Capture Application Manual*.

Statistics for monitoring the maximum hit size are not published in Web Console.

Restart PCA

To apply any of the above changes:

1. Save and close `ctc-conf.xml`.
2. Close any browser windows open to the PCA Web Console.
3. Restart the PCA. See "Installation" in the *IBM Tealeaf Passive Capture Application Manual*.

Configuring for the Pipeline

This section describes how to configure the maximum data size for pipeline hits.

Pipeline globals

You must configure the global setting for maximum data size that can be inserted into a Tealeaf pipeline (`MaxDataSize`) in the `[globals]` section of `TealeafCaptureSocket.cfg`.

- Child pipelines are not subject to these configuration limits, as they accept hits of any size.
- See "Overview of the Capture Pipeline and Session Agents" in the *IBM Tealeaf CX Configuration Manual*.

Inflate session agent

If the Inflate session agent is deployed in your Windows pipeline, you may encounter issues when hits are inflated to exceed the maximum size permitted by the session agent.

When you attempt to replay or search for hits that triggered this maximum size limit in the session agent, the following message may be displayed:

```
Page cannot be displayed. An error occurred while attempting to uncompress
the response for this page. The uncompressed data size exceeds the
configuration limit.
```

By default, the Inflate session agent limits the maximum inflated size (`Max Inflate Size`) to 512 KB. In this case, you must edit the Inflate session agent through the Pipeline Editor in TMS. Set the value to a higher value.

- See "Inflate Session Agent" in the *IBM Tealeaf CX Configuration Manual*.
- See Chapter 8, “TMS Pipeline Editor,” on page 307.

Configuring Session Size

You may also choose to review the maximum size settings for individual sessions captured in the Short Term Canister.

The Canister monitors the size of individual sessions as they grow in the Short-Term Canister across three dimensions:

- Hit count
- Data size in bytes
- Time in seconds

By default, Tealeaf configures a set of maximum limits for each metric and tests for them through a single event. This event can be changed through Javascript to test for other limits.

- See "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

Implementing Tealeaf in a Secure Environment

Tealeaf utilizes industry-standard conventions and mechanisms for providing secure capture, processing, and reporting of interactions with your web applications. Since your web servers may be capturing and sending sensitive data to your customers, Tealeaf provides both data- and access-level controls to ensure that sensitive information remains protected and secure. This section outlines the security features of the Tealeaf system.

- If you have additional questions, please contact http://support.tealeaf.com.

Pre-Install Considerations

Tealeaf administrators should be aware of any enterprise security policies that may be applied to any Tealeaf servers. Any one-time or periodic changes to user accounts, security policies, or other system configuration that may affect the performance or availability of Tealeaf software should be monitored by Tealeaf application administrators.

For example, some customers may reapply all their security settings as a nightly job to all machines, which may disrupt configurations or fixes applied to the Tealeaf system. These types of enterprise policies must be synchronized with the Tealeaf team.

PCA Security

The IBM Tealeaf CX Passive Capture Application is a passive network appliance that functions essentially as an advanced packet sniffer. The PCA maintains dedicated connections with the SPAN port or switch feeding data to it and should not have access to the wider Internet. As a result, security issues related to a web server having access to a live network connection do not apply.

- For secure communications, the PCA utilizes openssl, which is integrated into the application and is not patchable.

Operating System Security

The IBM Tealeaf CX Passive Capture Application is hosted on a separate Linux server, which can be secured by removing unneeded services and securing vulnerabilities. For more information, please contact your IT department or Tealeaf Professional Services.

Web Console

The PCA provides the Web Console, a web interface for configuring the application. The Web Console can be configured to transmit traffic over HTTPS.

Note: The Web Console is the PCA's administration console and is not intended for deployment over the open Internet.

- The PCA Web Console does not see any of the capture traffic and is dedicated to its web management functions.
- For more information on securing the Web Console, see "Hardware Setup and Operating System Installation" in the *IBM Tealeaf Passive Capture Application Manual*.
- For more information on the Web Console, see "Passive Capture Configuration via Web Console" in the *IBM Tealeaf Passive Capture Application Manual*.

If desired, the PCA Web Console can be disabled. Configuration of the IBM Tealeaf CX Passive Capture Application can then be applied through the configuration file (`ctc-conf.xml`).

- See "Passive Capture Configuration File ctc-conf.xml" in the *IBM Tealeaf Passive Capture Application Manual*.

If operating system changes are not possible or if editing the Web console configuration file is problematic, through the console you can specify specific traffic addresses and ports to ignore. Wildcards are accepted. See "PCA Web Console - Interface Tab" in the *IBM Tealeaf Passive Capture Application Manual*.

The Web Console is served by a standard Apache server. If desired, security enhancements can be applied through Apache.

- For more information, please see the Apache documentation.

Private Keys

This section describes how to configure your IBM Tealeaf servers for secure communications between the servers through SSL.

Applying Enterprise Private Keys to Tealeaf

To decrypt SSL communications, Tealeaf must be provided enterprise privacy keys. Tealeaf loads an encrypted version of the privacy key and decrypts it internally using a provided hash key.

Note: The PCA only handles SSL decryption where it handles SSL record processing discretely, which is likely to be different from an expected web server implementation and is thereby less vulnerable to attack.

1. Export the key in PEM format from any of the supported web servers. See "Exporting the SSL Private Key" in the *IBM Tealeaf Passive Capture Application Manual*.
2. Set up Tealeaf to use the exported key. See "Encrypted SSL Key Setup" in the *IBM Tealeaf Passive Capture Application Manual*.
3. You may also need to perform any of the following SSL key operations:
   - Generating a self-signed certificate
   - Generating a self-signed certificate using provided utility scripts
   - Setting up the Tealeaf Transport Service for SSL encryption
   - Setting up the PortalStatus/Web Console certificate
   - Remove or view certificate
   - Validating PEM keys
   - See "SSL Key Operations" in the *IBM Tealeaf Passive Capture Application Manual*.

Private keys can be added through SSL Keys tab on the Web Console, and you can review the keys that the PCA identifies as missing. See "PCA Web Console - SSL Keys Tab" in the *IBM Tealeaf Passive Capture Application Manual*.

Secure Communications with the Processing Server

The PCA can be configured to send secure communications (HTTPS/SSL) to the Tealeaf Processing Server or Servers in your environment. When enabled, the PCA delivers hits over HTTPS using a private key provided to the IBM Tealeaf CX Passive Capture Application.

Note: Enabling secure communications between the PCA and the Processing Server incurs an additional processing overhead for encryption and decryption operations, which impacts overall throughput.

Note: The PCA utilizes a special handshaking process with the receiving Windows-based components, which prevents transfer to any server other than a Tealeaf Processing Server.

See "SSL Key Operations" in the *IBM Tealeaf Passive Capture Application Manual*.

Integrating Tealeaf with HSM

Tealeaf can integrate with external Hardware Security Modules. See "Appendix - Integrating Tealeaf SSL Keys with HSM" in the *IBM Tealeaf Passive Capture Application Manual*.

Windows Server Security for IBM Tealeaf

This section describes some best practices for security with IBM Tealeaf and Windows-based servers.

Operating System Security

Any configuration settings or changes to the security of the operating system should be known in advance to Tealeaf administrators.

• See “Pre-Install Considerations” on page 356.

IIS Security

The Tealeaf Web Application Utility manages installation of the Tealeaf Portal and all file/group permissions required for IIS.

For enhanced security, you may consider purchasing and implementing an SSL certificate for the IIS web server to enable SSL access to the Tealeaf Portal.

Account permissions

For more information on the required account permissions, please review the following documents available from Microsoft:

- **IIS6:** http://support.microsoft.com/kb/812614/
- **IIS7:** http://support.microsoft.com/kb/981949/

Securing the Platform and Portal

This section describes how to use authentication and user accounts to configure a secure environment.

Portal Authentication Modes

The Tealeaf Portal supports multiple modes of authentication:

- NT Authentication
- Portal Authentication
- Single Sign On
- Mixed Mode (NT Authentication + Portal Authentication)

See “Authentication” on page 329.

To configure the authentication mode:

1. Login to the Tealeaf Portal as an administrator.
2. Select **Tealeaf** > **TMS**. The Tealeaf Management System is displayed.
   - See Chapter 4, "Tealeaf Management System," on page 259.
3. Click the Tealeaf node.
4. Select **Shared configuration information**. In the Config Actions pane, click **View/Edit**.
5. Click the Portal tab.

NT Authentication

Tealeaf can authenticate users using Active Directory through the NT domain with which it is associated. Tealeaf inherits all of the security features of NT Authentication.

- See "Authentication" on page 329.

Portal Authentication

The Portal provides its own authentication methods, in which user accounts are encrypted and stored in the Tealeaf database. Portal authentication enables enterprises to separate Tealeaf user management from their internal systems.

Portal authentication supports MD5, RC2, or 3DES encryption algorithms.

- See "Authentication" on page 329.

Mixed Mode

In Mixed Mode, users can choose to login through Portal Authentication or NT Authentication modes. The login screen provides a link to switch between these login modes, so that Tealeaf administrators can use the security features of both modes, as needed.

- See "Authentication" on page 329.

Single Sign On (SSO) Authentication

Tealeaf supports integration with CA Siteminder? or with other authentication providers that also use HTTP header requests. These logins behave like Portal Authentication logins. SSO authentication is configured via external file stored in the Tealeaf install directory.

- See "Authentication" on page 329.

Tealeaf Install Directory

On each server, the installed Tealeaf software components are stored in the Tealeaf install directory. This directory should be secured to prevent access by unauthorized local users and across the network. Only Tealeaf administrators should have read-write access to the Tealeaf install directory.

User Accounts

Depending on the enabled Tealeaf modules, one user account and one admin account are created by default. Through the Tealeaf Portal, Tealeaf administrators can configure password requirements and manage user accounts for IBM Tealeaf cxImpact, IBM Tealeaf cxReveal, and IBM Tealeaf cxView.

Password Security

For IBM Tealeaf cxImpact, IBM Tealeaf cxReveal, and IBM Tealeaf cxView, Tealeaf administrators can define the following password security options through the Portal Management page.

1. To access these settings, select **Tealeaf** > **Portal Management** in the Portal menu.
2. Select the appropriate settings to modify:
   - **IBM Tealeaf cxImpact** - Click the IBM Tealeaf CX Settings tab. Then, click **Users**.
   - **IBM Tealeaf cxView** - Click the IBM Tealeaf cxView Settings tab. Then, click **Users**.
   - **IBM Tealeaf cxReveal** - Click the IBM Tealeaf cxReveal Settings tab. Then, click **Users**.

Password Setting
Description

Minimum Password Length
Minimum password length for user passwords.
- Portal passwords can be a maximum of 32 characters in length.

Require Strong Passwords
Require passwords to meet the strong password requirements, which are combinations of upper-case, lower-case, and a number or a symbol.

Previous Password History (Count)
Number of unique passwords a user must have before reusing a password.

Previous Password History (Days)
Number of days that must pass before a user can reuse a password.

- See “CX User Administration” on page 165.
- See “cxView User Administration” on page 184.
- See "cxReveal User Administration" in the *IBM Tealeaf cxReveal Administration Manual*.

Additionally, you can specify whether accounts are automatically created for users when NT authentication is in use.

- For more information on IBM Tealeaf cxImpact user password settings, see “CX Settings” on page 192.
- For more information on IBM Tealeaf cxReveal user password settings, see "cxReveal Settings" in the *IBM Tealeaf cxReveal Administration Manual*.
- For more information on IBM Tealeaf cxView user password settings, see “cxView Settings” on page 216.

User Administration

For IBM Tealeaf cxImpact, IBM Tealeaf cxReveal, and IBM Tealeaf cxView, Tealeaf enables administration of users and user groups through the Tealeaf Portal. As needed at the user or group level, Tealeaf administrators can disable Portal menu features, force logouts after a time period, lock session replay mode, and other security-related features.

- See “CX User Administration” on page 165.
- See "cxReveal User Administration" in the *IBM Tealeaf cxReveal Administration Manual*.
- See “cxView User Administration” on page 184.

Tealeaf Canister Security

The Tealeaf Canister, or Processing Server, is used to assemble in-memory sessions that are currently active and then to store those sessions to disk when they are closed or otherwise time out. These sessions are then indexed for search.

- See "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

Tealeaf Databases

Depending on the installed components of your Tealeaf solution, the following database schemas need to be secured from unauthorized access.

Database Name
Tealeaf Component

TL_SYSTEM
Tealeaf system configuration database (introduced in Tealeaf 6.2).

TL_REPORTS
IBM Tealeaf cxImpact reporting database.

TL_RSEXTRACTOR
Segment Session analysis database.

TL_STATISTICS
Tealeaf Statistics database.

TL_VISSTAGE
IBM Tealeaf cxResults Staging database.

TL_VISREPORT
IBM Tealeaf cxResults Reporting database.

All applicable schemas should be secured from unauthorized access. Passwords to these databases should be distributed only to authorized personnel.

- By default the Tealeaf SQL database files are installed in the Tealeaf install directory.
- For more information about general database security, see "SQL Server Authentication" in the *IBM Tealeaf Databases Guide*.
- For more configuration options, see "Configuring the Report Server" in the *IBM Tealeaf CX Configuration Manual*.

Permissions

For more information on the required access and permissions, see "Database Manager SQL Server Access and Permission Requirements" in the *IBM Tealeaf Databases Guide*.

SQL Required Logins

Tealeaf requires two SQL logins (`TLADMIN` and `TLUSER`) to complete installation and management tasks for the Tealeaf databases.

A SQL login with sysadmin privileges is useful but not required for installation. If the SQL login does not have sysadmin privileges a DBA will be required to run scripts to complete the login. See "SQL Server Authentication" in the *IBM Tealeaf Databases Guide*.

RTV Security

RTV is a Windows application installed on the local desktop, which provides search and replay capabilities to Tealeaf users.

RTV makes requests to the Search Server to search and retrieve data. Depending on whether NT or Portal authentication is enabled, Search Server returns a different challenge string to which RTV responds in a way that's appropriate for the mode that it sees.

- For NT authentication, RTV uses the Windows account info of the current user.
- For Portal authentication, the username and password to use must be configured on the IBM Tealeaf cxImpact tab in the Options screen. See "RealiTea Viewer - Advanced Options Tabs" in the *IBM Tealeaf RealiTea Viewer User Manual*. RTV requests are made using the HTTP protocol. Data that is returned to RTV is always encrypted.

Best Practices

In general, Portal authentication and NT authentication provide equivalent data and access security. Depending on your specific enterprise needs, you may find the following sections helpful in making your choices.

Note: If NT authentication is used, the authentication master server must be a member of the domain to which the users belong, or the server must have sufficient rights to access group membership information for the users' domain. If that's not the case, Portal authentication is a better choice.

Data Security

For each of the following components, Tealeaf recommends the following security option for best results in securing data transmission across the network:

Tealeaf Component
Security Feature

IBM Tealeaf CX Passive Capture Application Server
Enable SSL connection.

IBM Tealeaf cxImpact Web Portal
Enable SSL on IIS

IIS Webserver
Enable SSL on IIS

Report server
NT authentication

Replay server
NT authentication

Unauthorized Login Access

For each of the following components, Tealeaf recommends the following security option for best results in securing the Tealeaf Portal from unauthorized access:

Tealeaf Component
Security Feature

IBM Tealeaf CX Passive Capture Application Server

- Enable NT authentication. See “Authentication” on page 329.

- Take steps to harden Linux. Please contact your IT department or Tealeaf Professional Services.

Report server
Enable NT authentication. See “Authentication” on page 329.

Replay server
Enable NT authentication. See “Authentication” on page 329.

Segment Builder Service

The Tealeaf Segment Builder service can be used to create session segments based on externally specified queries of Tealeaf session data. For example, a third-party application can specify a set of queries in a file, which are then executed to generate a new session segment.

Note: The Segment Builder Service is a component of IBM Tealeaf cxConnect for Web Analytics, a separately licensable component of the IBM Tealeaf CX system. please contact your IBM Tealeaf representative.

Uploading Segment queries

This service polls a preconfigured directory inside the Tealeaf install directory for query files, executes the queries, and assembles the returned session data into a session segment. The segment is then available through the Portal.

- When the license key permits use of the service, the Segment Builder Service is started by the Data Service automatically.

via Portal

The preferred method is to upload the file to the directory via the Tealeaf Portal. See "Managing Session Segments" in the *IBM Tealeaf cxResults User Manual*.

via File Directory

Prepared segment queries are retrieved from the following directory:

```
<Tealeaf install directory>\Integration\Segments
```

Operations

When the Segment Builder polls the `Segments` directory, it attempts to turn the `.txt`, `.csv`, or `.zip` files in the directory into sets of queries of the Tealeaf session dataset. If the process is successful:

1. The segment becomes available through the Managing Session Segments page. See "Managing Session Segments" in the *IBM Tealeaf cxResults User Manual*.
2. The source file is moved to the following directory:
   ```
   <Tealeaf install directory>\Integration\Completed
   ```
3. If there is a failure, the file is moved to the completed directory and the error is logged in the log file. In this case, no session segment is created.

Creating Segments for integration

The Segment Builder service can create segments out of source files in the following formats:

- `.txt`
- `.csv`
- `.zip`, which contains content in `.txt` or `.csv` format

Query options

If desired, you can specify the import settings for the queries in the session file. Otherwise, the default settings are applied.

Note: Unless specified in the import file, query parameters for the imported segments utilize the default parameter values as configured in the Portal. If you are regularly using the Segment Builder Service to import segments from a single external system, you should configure the Portal default values to support this integration, instead of manually configuring each import file to insert these values, as described below. For more information on the definitions of each setting and its default value, see "Managing Session Segments" in the *IBM Tealeaf cxResults User Manual*.

The following is an example query, as it would appear in a text file. In this example below, the search field that is queried is `ordernum` for the values `A123`, `B456`, or `C789`.

In the file, each query must be listed on a single, separate line:

```
#Options
#!MD5=N
#!DATA_COLUMN=0
#!DELIMITER=|
#!ID_PREFIX=NOTHING
#!SEGMENT_NAME=Order Number Search
#!DATE_SPAN=3
#!SEARCH_TERM=ordernum
#!SEGMENT_OWNER=admin
#!ID_PREFIX_USE=N
#!SHARE_SEGMENT=N
A123
B456
C789
```

Option Description

MD5 If set to `Y`, an MD5 hash is applied to the query string to normalize it for query execution.

DATA_COLUMN
The column in the data file that contains the search term values. If each line is a single value, then set this value to `0`.

DELIMITER
The column delimiter in the query file. The following values are supported:

- | (pipe)
- , (comma)
- (tab) - use keyword `TAB`

ID_PREFIX
The integration ID prefix that is used to identify search terms. If no integration ID prefix is used (`ID_PREFIX_USE=N`), set this value to `NOTHING`.

- For more information on this value, see "CX Settings" in the *IBM Tealeaf cxImpact Administration Manual*.

SEGMENT_NAME
The name of the segment as it appears in the Portal after a successful import.

DATE_SPAN
The number of days preceding the current date that are searched for matching data.

SEARCH_TERM
The variable name in the request data for which to search for matching values.

SEGMENT_OWNER
After import, this value identifies the owner of the segment among Tealeaf users.

ID_PREFIX_USE
Set this value to `Y` to use an integration ID prefix to mark the beginning of fields in the data.

- The value of the prefix is defined in the `ID_PREFIX` option above.
- The end of each field's data is defined by a delimiter (defined in the parameter `DELIMITER` above) or the end of the line.

SHARE_SEGMENT
To share this segment with all Tealeaf users after it is successfully imported, set this value to `Y`.

values You may specify values for which to search the field in `SEARCH_TERM` as individual lines for each value.

Query format

Each line in a segment file represents a different query of the Tealeaf database. The following options and settings dictate format:

Formatting Options
Description

MD5=N Data in text file for the search query is used. The MD5 hash is not applied before the search begins.

DATA_COLUMN=0
When this value is set to `0`, the resulting file is expected to be a single column of values.

DATE_SPAN=20
The query is created to search the preceding 20 days of data.

SEARCH_TERM
This value indicates the request variable field for which to search for values matching the column of data.

ID_PREFIX=N
No integration ID prefix is in use. The column of data just contains clear-text values.

The following options and settings are not relevant to this query:

- `DELIMITER`
- `ID_PREFIX=NOTHING`

After the query has executed and returned some sessions, the following settings are applied to the generated segment:

Option Description

SEGMENT_NAME
The user-readable name (`Testing only` in the above example) appears in the Managing Session Segments page.

SEGMENT_OWNER
The Tealeaf user ID of the owner of the generated segment: `joeuser`.

SHARE_SEGMENT
Since this value is set to `N`, the segment is a private segment visible only to `joeuser` and Tealeaf administrators.

Monitoring Segment Builder Service

You can monitor the progress of the Segment Builder Service through the Managing Session Segments page in the Tealeaf Portal.

The service detects any new files in the segment directory and processes them immediately. When the queries are finished, the assembled session segment is available through the Managing Session Segments page.

Note: For any imported segment in the Managing Session Segments page, verify that the Session Count column contains a non-zero value. In this column, a value of `0` indicates that the queries failed.

See "Managing Session Segments" in the *IBM Tealeaf cxResults User Manual*.

Configuring Segment Builder Service

All options for configuring the Segment Builder service are available through the Portal.

1. In the Portal menu, select **Tealeaf** > **Portal Management**.
2. In the left pane, select **IBM Tealeaf CX Settings**.
3. Select the Integration category.

See "CX Settings" in the *IBM Tealeaf cxImpact Administration Manual*.

Access permissions

For more information on configuring access permissions for segments, see "Report Access Permissions" in the *IBM Tealeaf cxImpact Administration Manual*.

Managing Your Tealeaf License Key

Beginning in Release 8.7, licenses for Tealeaf products are managed under the IBM Passport Advantage® system, and licenses are no longer embedded in the Tealeaf products.

Note: This section applies to Tealeaf customers who currently are using or upgrading to Release 8.6 or earlier or need to enable new products.

For Releases prior to Release 8.7, during the initial implementation and upgrades, Tealeaf provides each customer a new license key to enable access to all licensed components. Tealeaf components are tied to the existence and accessibility of a valid license key. Licensed components require a valid key to run and to enable functionality.

This section describes how to find and change the license key in a running system.

Required Upgrade Path

If you are upgrading to Release 8.0 or later, you must upgrade to Release 7.2.12.7296 before you upgrade to a Release 8 build.

Transitioning to IBM

As part of the transition to IBM, management of license keys for Tealeaf products is changing.

Release 8.7 or later

Note: For Release 8.7 or later, licenses to Tealeaf products are managed through the license issued by IBM to your enterprise.

- The Tealeaf license key is no longer used.

Release 7.2 or earlier

Note: If you have installed Release 7.2 or earlier, you cannot modify your license to enable new products. To enable new products, you must upgrade to Release 8.x or later. For more information, please contact Tealeaf Customer Support .

The remainder of this page applies to managing your license key for Release 8.6 or earlier.

Release 8.0 to Release 8.6

As needed, you may upgrade your Tealeaf solution or license new products.

The remainder of this page applies to managing your license key for Release 8.0 to Release 8.6.

IBM Passport Advantage

As part of the transition to IBM, license key authorization for existing Tealeaf customers has been moved into IBM Passport Advantage.

Release 8.0 to Release 8.6 customers must use the appropriate version of the License Key Generator tool available in IBM Passport Advantage to do the following:

1. regenerate your Tealeaf license key
2. modify your existing Tealeaf license key to enable new products

Please complete the following steps to update your license key.

Acquiring Your Current Key

The Tealeaf License Key should be viewed/updated through the Tealeaf Management System, where you can copy and paste the text file.

For steps to locate the key via TMS, see “Applying a License Key through the Portal” on page 372.

Acquiring the Tealeaf License File Generator

Enabling additional licensed Tealeaf components requires a new license key.

You must acquire a new license key if:

- You have licensed new Tealeaf products
- You are upgrading your Tealeaf system to Release 8.0 or Release 8.6
- You are directed to do so by Tealeaf Customer Support .

During upgrades, resynching the license key is especially important, as an out-of-sync key can result in lost functionality. If this loss is a database-driven component, the Tealeaf upgrader may not upgrade the proper database and the components.

Create an Account with IBM Passport Advantage

For more information on creating your IBM Passport Advantage account, see http://www-01.ibm.com/software/lotus/passportadvantage/.

Using the Tealeaf License File Generator

Through the Tealeaf License File Generator utility, you can enable all products and components that you have licensed.

- Based on your selections, however, some options are not available.

Expiration

In the upper-right corner, you can configure the expiration date for the license key. The current expiration date is highlighted in blue.

- By default, the utility configures the license key to expire fourteen days after the current date.
  - Expiration time is based off of GMT, since the utility has no knowledge of the Tealeaf system timezone for the affected system.
- To change the expiration date, use the calendar controls and select a new date.
- To create a perpetual license, click the No Expiration checkbox. For customer production environments, perpetual licenses are used.

Portal Modules

Table 39. Platform:

Module Name	Description
`cxImpact`	The Standard Edition of the Tealeaf platform. **Note:** To access the Tealeaf Portal, the IBM Tealeaf CX solution must have a IBM Tealeaf cxImpact license, Standard or Extended.
`Browser Replay`	Through the Portal, Tealeaf users can replay captured and archived Tealeaf sessions.
`cxImpact - Extended`	(post-7.2.8.7260 version) Enables the Extended Edition of the Tealeaf platform, which includes UI Capture SDK and related Portal reports. **Note:** To access the Tealeaf Portal, the IBM Tealeaf CX solution must have a IBM Tealeaf cxImpact license, Standard or Extended.
`UI Capture`	Enables capture of client user interface events.

Table 40. Applications:

Module Name	Description
cxView	IBM Tealeaf cxView enables scorecard and dashboard functionality in the Tealeaf Portal.
Scorecards	Scorecards enable report snapshots in table, chart, or graph form, which can be emailed to stakeholders as needed.
Dashboards	Through dashboards, individual reports and scorecards can be collected into a Portal-based dashboard that can be configured by the individual Tealeaf user.
cxReveal	Tealeaf solution for the Call Center.
Database Search	When enabled, you are licensed to install and deploy the IBM Tealeaf cxReveal Search database, which enables capture and searching for sessions by session attribute (not available for Release 8.0). • Software installation is required. • Beginning in Release 8.3, Database Search is automatically enabled when IBM Tealeaf cxReveal is licensed.
cxResults	IBM Tealeaf cxResults enables enhanced reporting on customer experience and key performance indicators. **Note:** As of Release 8.7, IBM Tealeaf cxResults is withdrawn from market and is no longer available for new Tealeaf installations.
Session Segments	Session segments are sets of captured sessions and fragments of sessions that can be managed and analyzed through the Tealeaf Portal. Selecting this option enables all RSE Service functionality.
Visitor Segments	Visitor segments are sets of sessions captured by individual visitor. Visitor segments can be managed and analyzed through the Tealeaf Portal.
cxConnect	IBM Tealeaf cxConnect for Data Analysis enables export of Tealeaf session data for integration with third-party platforms. • For new licensees of IBM Tealeaf cxConnect for Data Analysis, additional software installation is required.
cxVerify	IBM Tealeaf cxVerify provides packaging and storage for Tealeaf archives to the enterprise data warehouse. **Note:** As of Release 8.7, IBM Tealeaf cxVerify is withdrawn from market and is no longer available for new Tealeaf installations.
cxMobile	Enables the IBM Tealeaf CX Mobile module, which provides tools and events for capturing mobile user activities.
cxOverstat	Enables the IBM Tealeaf cxOverstat product, which enables the capture of usability data from visitor sessions and display through BBR. This product was introduced in Release 8.6. **Note:** IBM Tealeaf cxOverstat requires the installation of the JSON version of the UI Capture SDK, which is a component of the IBM Tealeaf CX-Extended license. IBM Tealeaf cxOverstat requires this license as part of its license key.
Javascript Advanced Mode	In Release 8.0 to Release 8.2, this option enables access to the Javascript editor through the Tealeaf Event Manager. **Note:** As of Release 8.3.0.8330, this options is automatically granted to all licensees.

Licensing Dependencies

The following licensing dependencies are reflected in the options that are available for selection:

Platform Modules

Module
: **Rules**

cxImpact
: Standard Edition of the IBM Tealeaf CX platform

cxImpact - Extended
: Extended Edition of the IBM Tealeaf CX platform, which enables UI Capture SDK and related reports.

  - During the upgrade process from Release 7.2.8.7260 or earlier, existing Tealeaf customers are granted the IBM Tealeaf cxImpact - Extended license automatically.

Feature and Product Modules

Module
: **Rules**

Browser Based Replay
: If you license either IBM Tealeaf cxImpact or IBM Tealeaf cxReveal, Browser Based Replay is included automatically.

UI Capture
: UI Capture is included with the IBM Tealeaf cxImpact - Extended license.

Scorecards and Dashboards
: Included with the IBM Tealeaf cxView license, scorecards and dashboards can be enabled or disabled separately, too.

cxResults
: IBM Tealeaf cxResults provides the ability to search, analyze, and report on visitor sessions, including the use of `Session Segments` and `Visitor Segments`.

  - `Visitors Segments` is available only through IBM Tealeaf cxResults. If you license one, you get the other.
  - `Session Segments` were previously included in the IBM Tealeaf cxImpact license. For Release 7.0 or later, they can be enabled separately.

    Note: During the upgrade process from pre-Release 7.0 systems to Release 7.0 or later, the ability to analyze session segments is "grandfathered" into the new installation, regardless of IBM Tealeaf cxResults licensing.

cxVerify and cxConnect
: third-party integration modules that can be licensed independent of any other licensing.

cxMobile
: Enables cxMobile Experience Management module capabilities; requires the IBM Tealeaf cxImpact or IBM Tealeaf cxImpact - Extended license.

cxOverstat
: Enables IBM Tealeaf cxOverstat product capabilities

Advanced Modules

Note: These licensing options are for internal Tealeaf testing purposes only. None of these options should be enabled in a customer installation.

Module
Description

Static Demo Data
This option enables reporting data feeds from a demo data set that has been previously installed on the local system.

Note: This setting should not be enabled by Tealeaf customers. Some Portal features are broken when this option is enabled.

Dashboard Custom SQL
This option enables some super administrator features in the Portal.

Note: This setting should not be enabled by Tealeaf customers. Some Portal features are broken when this option is enabled.

Customer Information

Note: This section applies to Release 8.0 to Release 8.6 only.

A customer's license is unique to the customer. Enter the NameSpace and Name, as defined below.

Property
Description

NameSpace
Enter the namespace in the space provided. This value should be provided to you by Tealeaf Customer Support .

Note:

1. Do not begin a customer name space value with a numeric value (0-9). Numeric values cannot be used at the beginning of the NameSpace value.
2. Do not use the value TL, which is reserved for Tealeaf.

Generating a License Key

Note: Do not attempt to generate a key for a Tealeaf release version other than the version from which you are using the utility. Before you generate and apply a license key, please verify that build version of the License File Generator utility corresponds to the version of the Tealeaf Portal that you have installed or to which you are upgrading. For existing Portal installations, the build number is listed at the bottom of each page in the Portal in the following form: A.B.C.DDDD.

For Upgrades

Steps: If you are upgrading or enabling new features, please complete the following steps first:

1. Acquire the license key through TMS.
   - For steps to locate the key via TMS, see "Applying a License Key through the Portal" on page 372.
2. Acquire and launch the version of the Tealeaf License File Generator for the build to which you are upgrading.
   - For more information, see "Acquiring the Tealeaf License File Generator" on page 368.

3. Paste the key into the Tealeaf License File Generator.
4. Click **Validate**. For more information, see "Validating a License Key."
5. Verify that the current configuration is reflected in the selected components.
6. Continue with the following steps, selecting any new modules as needed.

For New Installs

1. Acquire and launch the version of the Tealeaf License File Generator for the build to which you are installing.
2. Select the Portal modules to include.

   Note: Based on your selections, some modules may not be available.
3. Select any Advanced modules that apply.

   Note: Be careful when enabling Advanced modules. Some are intended for debugging purposes only.
4. Enter the NameSpace and Name values provided to you by Tealeaf Customer Support .
5. Set the expiration date in the calendar.
6. To generate the key, click **Generate Key**.
7. This key must be provided to the Tealeaf installer/upgrader, when prompted.
   a. **Old Version:**
      - Copy and paste the key into the appropriate textbox.
   b. **post-Release 7.2.8.7260 version:**
      - To copy the key to the clipboard, click **Copy to Clipboard**. In the installer/upgrader, paste the key into the appropriate textbox.
      - To export the key into a Tealeaf license file, click **Save License File**.

Validating a License Key

You can use the utility to test a license key that has been provided to you.

To validate a license key:

1. Enter or copy and paste the license key into the License textbox.
2. Click **Validate Key**.
3. The following message should appear:

   ```
   License key is valid.
   ```

Applying a License Key through the Portal

Note: If your new license key includes enabling new Tealeaf products that require the installation of new databases, you cannot install the license key through the Portal. The new license key must be applied through the Tealeaf Installer as part of the installation of the new products. These products include:

- IBM TealeafcxReveal with Database Search (build 7.2.12.7296 or later)
- IBM TealeafcxResults (Release 7.0 or later)
- For more information, please contact Tealeaf Customer Support .

After you have acquired a new or updated license key, you can apply it to your Tealeaf solution through the Tealeaf Management System, which is accessible through the Tealeaf Portal.

1. In the License File Generator utility, generate and validate a new license key.
2. Click **Copy to Clipboard**.
3. Log into the Tealeaf Portal using an account that has access to the Tealeaf Management System.
4. From the Portal menu, select **Tealeaf** > **TMS**.
   - If you cannot see this menu option, your account does not have access to TMS. For more information, please contact your Tealeaf administrator.
5. In TMS, select the Servers view.
6. Expand the node for the (Master) system and then expand the `Tealeaf` node.
7. Select `Global configuration information`.
8. Click **View/Edit (Raw)** in the Config Info pane.
9. The Tealeaf global configuration settings window is displayed. Click the Show Hidden Items checkbox.

   Note: Typically, you should avoid editing hidden configuration items. These items are removed from public view, as changing them can have significant impacts on system performance if they are mis-configured.
10. Click the License value.
11. Paste the license key into the textbox, including the start and end tags, as indicated below:

    ```
    -----BEGIN TEALEAF LICENSE-----
    (license data)
    -----END TEALEAF LICENSE-----%00
    ```
12. Click **Apply**.
13. Click **Save**.
14. The changed license key must be applied to any other Tealeaf servers in your solution. To apply the license key immediately, click **Add Tasks and Submit**.
15. The task to update the license key is submitted, and the other servers are updated.
    - See “TMS WorldView Tab” on page 262.

Applying the License Key through the Registry

If your license has expired, the Portal and the Tealeaf Management System may no longer be available. If so, an updated license key can be applied through the Windows Registry.

Note: Whenever possible, apply the Tealeaf license key through TMS.

- See "Troubleshooting - Portal" in the *IBM Tealeaf Troubleshooting Guide*.

When Your License Key Expires

When your Tealeaf license key expires, the following systems no longer work:

- Tealeaf users cannot access the Portal. The Portal API is no longer available.
  - Since it is accessed through the Portal, Browser-Based Replay is not available.
  - Users can continue to search for and replay sessions through RTV.

- The Report Server stops working. Data collection stops. The databases are no longer updated, and the Portal-based reports that draw from them receive no new data.

The Tealeaf PCA continues to capture data, process it, and forward it to the Processing Servers, which continue to evaluate incoming hits. However, since the Data Collector is no longer operational, data in the Canister begins to expire within 3 days. This data cannot be recovered.

Note: If your license key expires, you have three days to renew the license before data loss begins.

Managing Static Archives

The IBM Tealeaf CX system is designed to capture the minimum volume of data that is required for high-fidelity replay and detailed reporting on web application activities. Much of the data that is experienced during a visitor session is repeated from session to session. Multiple visitors may request the same images, JavaScripts, or style sheets. Capturing each instance of these data objects is expensive for storage and may impact processing throughput. Therefore, TLI does not scale to capture all static content.

By default, at the point of capture, the IBM Tealeaf CX system is designed to drop this content, and during replay, the Tealeaf client applications refer to the origin server to download this static content. This design minimizes the data that must be captured and processed for enable high-fidelity replay.

Note: Because of limitations in using TLI to capture all static content, IBM Professional Services should review your deployment to ensure successful deployment of scale.

In some environments, this solution is not optimal for the following reasons:

- Contacting the origin server during the replay may violate enterprise security policies.
- Direct access to the origin server may not be possible without authentication, may be blocked by a proxy server, or may be masked by a content delivery network.
- Since the static content on the origin server may change over time, references within session data to static content on the origin server may fail to retrieve content or may pull back the wrong versions of content.
- For legal purposes, a snapshot of a visitor's session may be a requirement.
- Touching the origin server during replay may increment usage statistics that are tracked by third-party systems, which corrupts true measurement of site activity and performance.

For environments impacted by one or more of the above conditions, Tealeaf supports the automated creation of static content databases that are populated based on the capture stream that is passed through the Windows pipeline. When properly enabled and configured, the TLI session agent scans the data passing through the pipeline for a configured set of mimetypes and inserts those data objects into a static object database.

- This static object database is rolled on a daily basis, creating an effective archive of all of the static content experienced in visitor sessions for that day.

- Static archives can be made centrally available for replay and stored for data warehousing purposes.

This section describes how to configure and enable the automated creation and maintenance of Tealeaf static archives.

Note: Users of the IBM Tealeaf CX RealiTea Viewer, the standalone desktop application for session replay, can create local static archives for storing content. This content is available only to the user of the desktop system where it is stored.

Overview of Static Archives

A **static archive** is a database of static content assets extracted from session data. This database is stored in a proprietary file format (`.TLI`).

Each static archive represents a snapshot of static content detected in session data. Depending on the type of static archive, this snapshot may represent a day, a month, or (for local archives) a manually determined period.

The following types of static archives are available:

- **"Local Archives"** - Through the IBM Tealeaf CX RealiTea Viewer, end users can create static archives that are stored and maintained on their local machines. These archives must be manually updated and are not automatically shared with other users. See "Local Archives."
- **"TLI Server Archives" on page 376** - Tealeaf administrators can install and configure a TLI server, which automatically creates and maintains daily and monthly static archives from content extracted in the Windows pipeline. See "TLI Server Archives" on page 376.

Local Archives

Through RTV, Tealeaf users can manage their own static archives. In addition to saving session data to the local hard drive, RTV users can store the accompanying static content in a `.TLI` file, which creates a completely self-contained archive.

- Local static archives are updated based upon the references in a currently open session. When the update is executed, the static references in the session data are checked on the origin server and, if changes have occurred, updated into the local static archive.

Storing local instances of static content simply distributes the data storage requirements across multiple desktops, and most of the content stored by individuals is duplicated on other desktops. From an enterprise storage perspective, the use of local static archives can be expensive.

Note: In some web applications, static content is loaded based on the dynamic execution of Javascript. Execution of Javascript may be disabled in RTV for a variety of reasons. When Javascript is disabled, you may not be able to retrieve a full set of static content for the current session.

- In the Static Files Options tab, RTV users can define their own static archives.
- Using the GetImages command, the active static archive can be updated at any time based on the referenced in the currently open session.
- As needed, you can view the content in your local static archive.
- See "Using Static Archives in RTV" in the *IBM Tealeaf RealiTea Viewer User Manual*.

Note: The remainder of this section describes how to manage server-side static archives.

TLI Server Archives

Tealeaf provides a mechanism for automated creation and updating of static archives through the Windows pipeline. When the pipeline and required session agent are configured to capture static content types, static archives are created and rolled on a daily basis.

During replay of sessions, Search Server is instructed to scan the appropriate static archive based on the timestamps in the session data. For example, if the timestamp for a hit indicates a date of December 14, 2010, Search Server looks for static content in the corresponding static archive for that date. If it fails to find the appropriate content, it may examine archives for adjoining dates for the content.

- Search Server works backward from timestamp of the hit to find the first unique instance of that object.

Architecture

This section provides an overview of how static content is stored, used at runtime, impacts storage and throughput, and lives in the Tealeaf system.

How Static Content is Stored

A server-side static archive created by the TLI session agent is stored in `.TLI` format. This Tealeaf format is a ctree file format, which can be configured to store any type of data object. Embedded in the file format are the indexes used to search the file. This all-in-one solution enables easy management of static archives.

The ctree disk-based database is used in the Tealeaf solution for storage of completed sessions in the Long-Term Canister. Like the LTC, the `.TLI` file can be searched using the Tealeaf Search Server interface, which is queried during replay for the appropriate `.TLI` file to scan for the objects to replay.

For each data object stored in the `.TLI` file, a unique identifier is generated using the URL and an MD5 checksum applied to the content, which enables the session agent to later detect variations in newly detected instances of the same object.

- Prior to calculating the checksum, headers are stripped from the object, as they may contain variable data.
- The combination of URL and checksum provides a unique key for identifying and locating static content in the archives.

When new instances of the same object are detected, the checksums of the two objects are compared.

- If the URLs and checksums differ, then the second object is stored in the TLI database as a separate object.
- If the URLs differ and the checksums are the same, then the second object is recorded as a reference to the first object.

Note: Response types are inserted only if the HTTP Status Code is 200 (ok) and if content length is greater than zero.

Note: Insertion of static content into the TLI archive is dependent on resolving the mimetype of the content and configuring the PCA and the Windows pipeline session agent to recognize that mimetype. However, some static content, such as

Status Code 302 responses, does not have a mimetype assigned to it. If you have questions about resolving these issues, please contact Tealeaf Professional Services.

How Static Archives Are Used at Runtime

During replay, the timestamps recorded for the hit are used to make the best match of content. If there are multiple instances of the same object, the match is performed based on date/time information.

- During replay through the IBM Tealeaf CX RealiTea Viewer, RTV first checks to see if the static object has been stored in its local cache. If not, it then submits a request to Search Server to find it.

Search Server attempts to search the static archive from the matching date for the matching object. If no match is found, Search Server attempts to search the archives from adjacent dates for the object.

- If a match is found, RTV caches the object in a set of local TLI cache files. For this reason, you may wish to clear your IE Cache in RTV from time to time.
- You can identify the objects that have been loaded from the local TLI cache through RTV. See "RealiTea Viewer - Page Load Details" in the *IBM Tealeaf RealiTea Viewer User Manual*.

Assess Content Types to Capture

The capture of static content is based upon the content type specified in the response of the data.

Typical static content types include the following. You should create your own list:

Content Type
Description

image/gif
GIF image

image/jpeg
JPEG image

image/png
PNG image

image/x-icon
icon

text/css
Cascading stylesheet

text/javascript
Javascript

application/javascript
Javascript application

application/x-javascript
Javascript application

In addition to the above, your site may contain other static content types for PDF files, movies, and other high-density formats.

Note: Depending on the size of the average response and volume of responses of these types, you may wish to enable the content type with the lowest average daily

size first, in order to test TLI functionality. However, you should verify that the content type appears frequently enough to be captured in a reasonable sample of sessions.

Assess Impact on Data Storage and Throughput

Note: Static content, such as image files or PDFs, can significantly increase the volume of data processed by Tealeaf, which can thereby impact overall throughput in the system. Please be sure to review this section for assistance in gauging storage impact. For more information, please contact Tealeaf Professional Services.

Data storage impact

There are no definitive tests or measurements for assessing the impact of capturing static content. You may be able to estimate the size in terms of daily storage requirements based on static content that you capture locally through RTV.

To perform this capture, please complete the following steps:

Note: Where possible and appropriate, the TLI session agent compresses data prior to storage.

1. Create a representative set of sessions that can be used to estimate the volume of unique instances of static data served on a daily basis. Browse your web application to create the session or sessions.

   Note: Since static archives are designed to collect only one copy of each instance of static content, you can create sample sessions of any size and any number to determine your sample static content. The important thing is to gather a statistically significant volume of sessions and data.

   a. Create sample sessions in a way that makes them easy to find through search. For example, if you use a consistent login identifier, you can search for that data to find your sessions.
   b. For smaller sites, you may be able to create a session in which you browse to every page on the site. This method creates a session containing references to all static content.
   c. For larger sites, you might gather a representative set of sessions from a single day and then use RTV to gather the static content from these sessions into a single static archive. For example, you could gather a sample of five medium-sized sessions and use that as your baseline metric for the size of each day's static content archive.
2. Through RTV, you can create a static archive to store the content.
   a. In the RTV menu, select **Tools** > **Options...**.
   b. Click the Static Files tab.
   c. Create the new static archive and move it to the top of the list. The topmost archive is marked with an asterisk to indicate that it is the currently active one.
   d. Select the GetImages checkbox.
3. Then, you open your session or sessions and perform a GetImages command to store the static content referenced in each session into the static archive.
   a. To perform a GetImages, select **Tools** > **GetImages**.
   b. Static content referenced in the session data is downloaded and stored in the static archive you just created.
4. Find the size of the static archive.

a. In the RTV menu, select **View** > **Static File Database...**.
b. In the dialog, the contents of the static archive are displayed.
c. The volume of data is indicated by the Total Bytes metric at the top of the window.
d. For further analysis, you can export the static file information to Excel. If you sub-total the data by ContentType, you may be able to determine the mimetypes that will have the greatest impact on storage and throughput. These items are candidates for enabling last.

See "Using Static Archives in RTV" in the *IBM Tealeaf RealiTea Viewer User Manual*.

Throughput impact

Before you begin deploying static archives, you should assess a baseline performance of your Tealeaf pipelines. You can use the Pipeline Status tool to track the rate at which your main Windows pipeline or pipelines process session data.

This baseline becomes the basis for comparison after you have deployed static archives. When you have completed the configuration and begin capturing static content to archives, you can use the Pipeline Status tab to assess the rate of insertion into the static archives and the impact elsewhere in the Windows processing pipeline.

- See “Monitoring Static Archive Pipeline Status” on page 388.

Deployment Models

While it is possible for the TLI session agent to be deployed in the same main processing pipeline, Tealeaf recommends that you create a separate pipeline that is responsible for capturing static content only. For best results, static archives should be managed on a dedicated machine in a single Windows pipeline whose sole purpose is to manage static archives.

Tealeaf recommends that you deploy the RTA Split session agent to deliver the static content to the machine and/or pipeline where the content is to be inserted into the TLI archive.

- The RTA Split session agent should be inserted into the main processing pipeline as early as possible to limit the amount of unnecessary or duplicate processing within the pipeline(s).
  - See "TLI Session Agent" in the *IBM Tealeaf CX Configuration Manual*.
  - * See "RTA Split Session Agent" in the *IBM Tealeaf CX Configuration Manual*.
- Depending on the type of static content that is being captured and your enterprise security requirements, privacy rules may need to be applied to it before RTA Split delivers it to the static content processing pipeline. In these cases, RTA Split should be inserted after the Privacy session agent. See "Privacy Session Agent" in the *IBM Tealeaf CX Configuration Manual*.
- For more information on applying the session agents and configuring them, see Chapter 8, “TMS Pipeline Editor,” on page 307.

HBR environments

Health-Based Routing can be deployed in Tealeaf environments to manage the distribution of Tealeaf hits between multiple Canisters.

Note: Health-Based Routing requires significant system resources. The system hosting HBR should not be burdened with static content processing. If you are deploying static archiving into a Tealeaf environment, do not deploy onto the HBR machine itself.

Note: If you are using Health-Based Routing in your Tealeaf environment, please consult with Tealeaf Professional Services to determine the best approach to managing the static content traffic.

See "Health-Based Routing (HBR) Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

If you have deployed HBR, you can use the RTA Split session agent to route traffic to the dedicated TLI server for capture through the TLI session agent.

- Instructions on how to enable this configuration are described later.

As a separate Canister pipeline

If the above deployment is not possible, Tealeaf recommends deploying a separate pipeline for capturing static content on a server that is not otherwise heavily utilized.

Note: In the PCA, do not clone all traffic for delivery between the main Windows pipeline processing system and the static content machine. Cloning results in a large and inefficient increase in content delivered over the network.

Configuration

Note: Creation and management of server-side static archives requires that IBM Tealeaf cxImpact Build 7.2.12.7296 or later. Use of these archives in replay through RTV has the same requirement. Please upgrade to the latest available versions. For more information on downloading IBM Tealeaf, see IBM Passport Advantage Online.

To enable centralized management of static archives, the following configuration steps need to be executed.

1. Install the necessary software on the TLI server. See “1. Install Tealeaf Software” on page 381.
2. You must create references to the TLI Server in the Portal, which enables the Portal to gather status information on static archive activities. See “2. Configure a TLI Server Reference” on page 381.
3. Determine the configuration of the Windows pipeline or pipelines designated for capture of static content. See “3. Configure the Static Archive Pipeline” on page 382.
4. The TLI session agent must be deployed and configured in the designated pipeline. See “3. Configure the Static Archive Pipeline” on page 382.
5. By default, the IBM Tealeaf CX Passive Capture Application is configured to drop the most common forms of static content. Through the PCA Web Console, you must enable the capture of static content mimetypes.
   - See “4. Configure PCA Capture Types” on page 382.

     Note: After you have completed the changes in the PCA Web Console, you must restart capture. See "Installation" in the *IBM Tealeaf Passive Capture Application Manual*.

6. Configure your Replay Servers to use the TLI server. See “5. Configure Replay Server Use of TLI Servers” on page 383.
7. After the above steps are completed, please allow a period of at least one hour to pass during which the first static archive is populated with content.
8. When the static archive has been populated, you can verify operations through IBM Tealeaf CX RealiTea Viewer. See “Testing Static Archives” on page 383.
9. You can monitor the changes in throughput rates through the TMS Pipeline Status tab. See “Throughput impact” on page 379.
10. To manage static archives on the TLI server, you should configure merge and trim operations through the Scheduling Service. See “Managing TLI Files” on page 389.

1. Install Tealeaf Software

Note: Tealeaf recommends processing static content through a dedicated Windows pipeline on a standalone server.

If you are installing the TLI server on a standalone system, you must use the IBM Tealeaf cxImpact Installer to install the Transport Service without HBR on the system.

- See "CX Installation and Setup" in the *IBM Tealeaf CX Installation Manual*.

Install and Configure RTV

Before you begin, you should verify that you have installed the latest available version of the IBM Tealeaf CX RealiTea Viewer desktop application, which can be used to verify replay of static content and to check the contents of TLI files.

- See "RealiTea Viewer (RTV) User Manual" in the *IBM Tealeaf RealiTea Viewer User Manual*.
- After you have installed the application, you must configure it to work with static content. See "Using Static Archives in RTV" in the *IBM Tealeaf RealiTea Viewer User Manual*.

2. Configure a TLI Server Reference

After you have installed the Transport Service on the designated TLI server, you must create a reference to the TLI Server through the Portal Management page. Adding this configuration entity enables the Portal application to contact the TLI server and to extract status information, including the current TLI files on the server.

- To access the Portal Management page, select **Tealeaf** > **Portal Management** in the Portal menu.
- Additionally, this entry enables the Portal to find the TLI Servers from which to extract sessions for the data warehouse.

  Note: The TLI Server and Search Server must use the same port. By default, this value is configured to `19000`.
- See “Managing Tealeaf Servers” on page 115.
- For more information on reporting, see “Static Archives in System Status” on page 388.

3. Configure the Static Archive Pipeline

After you have installed and deployed the TLI server, you should create a dedicated Windows pipeline to capture static content to `.TLI` files.

This configuration involves:

- **Configure the pipeline to capture the static content:** Configured on the designated TLI server, this pipeline must include the TLI session agent.
- **Deploy the RTA Split session agent to deliver the static content:** When the TLI server and pipeline have been configured, you must deploy the RTA Split session agent to send static content to the TLI server's pipeline. This session agent must be deployed in any pipeline that processes hits, either at the Canister or HBR tier.

For more information on configuring these items, see "TLI Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

4. Configure PCA Capture Types

Capture types that you have configured the TLI session agent to insert into the static archives must also be enabled for capture through the IBM Tealeaf CX Passive Capture Application.

In most implementations, the PCA is configured to drop from capture binary and other static content. Through the PCA Web Console, you can enable the capture of these types for processing through the system.

Note: Enable the PCA to capture static content as the last step of the configuration. When the PCA begins capturing the content, the remainder of the static archive system needs to be able to handle the sudden increase in traffic. You may wish to enable capture of only one static content type to see the effects on capture volume first.

Depending on the static content of your web application, you may need to modify this list.

Note: During deployment of static archives, you might choose to initially enable the capture of a single mimetype that typically does not reference large files, which may significantly impact the system. After you have enabled capture of a type such as `image/gif`, you can verify static file operations. Using the same change-then-test iteration process, you can enable capture of all required mimetypes while monitoring the impacts of each type on the system.

Enable for capture

- **Mimetypes:** The static content mimetypes must be enabled for capture through the Web Console of the IBM Tealeaf CX Passive Capture Application.
  - See "PCA Web Console - Pipeline Tab" in the *IBM Tealeaf Passive Capture Application Manual*.
- **File extensions:** The file extensions must be enabled for capture, as well, through the PCA.
  - Verify that the file extensions are not excluded. See "PCA Web Console - Pipeline Tab" in the *IBM Tealeaf Passive Capture Application Manual*.
  - Verify that the file extensions are included. See "PCA Web Console - Pipeline Tab" in the *IBM Tealeaf Passive Capture Application Manual*.

5. Configure Replay Server Use of TLI Servers

After you have created a TLI server and enabled the capture of static content throughout the Tealeaf system, you must configure Replay Server to use the TLI server to retrieve static objects. Through Replay Server, Browser-Based Replay users can access static content during replay through the browser.

1. Login to the Portal as an administrator.
2. From the Portal menu, select **Tealeaf** > **TMS**.
3. Click the WorldView tab.
4. From the View drop-down, select `Servers`.
5. Select the server hosting the Replay Server.
6. Open the Replay Server node.
7. Click **Replay Server configuration**.
8. Click **View/Edit**.
9. The Replay Server configuration is displayed. The following settings need to be configured:

   Setting Description

   `Enable TLI Server`
   Set to 1 to enable use of a TLI server by the Replay Server.

   `TLI Cache Path`
   Configure the full path to the directory on the Replay Server where the static content cache is maintained.

   Note: Objects in the cache are expired after one day.

   `TLI Server`
   Set this value to the name of the machine hosting the TLI server.

When the above items have been configured, click **Save**.

Note: The above configuration change needs to be applied to all Replay Servers in your environment.

Testing Static Archives

As soon as you have configured static archiving and enable the capture of one static content mimetype through the IBM Tealeaf CX Passive Capture Application, you can begin testing the capture of static content to a TLI file through the IBM Tealeaf CX RealiTea Viewer.

Note: Since a static archive is stored as a file on the hosting server, it can be queried through RTV for real-time updates. As soon as the content is written to the archive, it is available for review through RTV.

Steps: In this example testing procedure, the enabled capture type is `image/gif` for `GIF` image files.

- If you have not done so already, install the latest version of RTV. See "RealiTea Viewer (RTV) User Manual" in the *IBM Tealeaf RealiTea Viewer User Manual*.
- You must also configure RTV to work with static archives. See "Using Static Archives in RTV" in the *IBM Tealeaf RealiTea Viewer User Manual*.
- Generate a session in your web application that accesses a page referencing at least one instance of a `.GIF` file.

- Since the hits are processed in real-time, it should be a matter of a few seconds before the `.GIF` file has been inserted into the TLI file.
- If you have properly configured RTV, you should be able to see the static content immediately after it has been inserted into the TLI file. From the RTV menu, select **View** > **Static File Database...**.
- The contents of the currently active static archive on the selected TLI server are displayed. In the list, locate the `GIF` file that you visited.
  1. If you find the reference, check the Timestamp column to verify that the entry closely approximates when you visited the page. If it does not, there may be a processing issue in the pipeline.
  2. If you do not find the reference but can see other GIF files, the hit may not have yet been processed.
     a. You may view the static database at a later time.
     b. Be sure to check the timestamps of other GIF files to verify that they have been captured after you enabled the capture type in the PCA. If the timestamps are from before the time when you enabled capture, you should clear your IE Cache of objects and re-verify your RTV configuration. See "Using Static Archives in RTV" in the *IBM Tealeaf RealiTea Viewer User Manual*.
  3. If the Static File Database Contents dialog is empty, you may be using the wrong type of TLI capture method or may be having another configuration issue in RTV. See "Using Static Archives in RTV" in the *IBM Tealeaf RealiTea Viewer User Manual*.
- If you are still experiencing issues with static content capture, you should verify operations of Tealeaf processing through Portal-based reporting. See “Static Archives in System Status” on page 388.
- You may repeat the above process after enabling other capture types through the PCA. See “4. Configure PCA Capture Types” on page 382.

Actions on Static Archives

This section describes the actions that can be used to manage your archives.

Static Archives Lifecycle

Over time, the static content on your web application is likely to change. Management of static archives must adapt to these changes while preserving high-fidelity capture, replay, and warehousing of sessions.

To maintain adaptability while preserving data integrity, Tealeaf supports the following functions for static archives:

- **Daily Roll** - Each day, the active static archive is rolled. The current TLI file is closed and saved, and a new one is created to store static content detected in the capture stream during the coming day. See “Rolling Archives” on page 385.
  – You may wish to backup the TLI files on a daily basis using your enterprise backup procedures. See "TLI Session Agent" in the *IBM Tealeaf CX Configuration Manual*.
- **TLI Merge** - Each day, after the daily static archive has been rolled, you can configure a task to roll the daily static archive into a monthly file for long-term record keeping and optionally move it in the TLI backup directory. See "Configuring the Scheduling Service" in the *IBM Tealeaf CX Configuration Manual*.
- **Trim** - Static archiving supports trim operations applied to all TLI files stored on a TLI server. You can configure this operation to occur on a daily basis to

remove static content with timestamps older than a specified number of days. This feature is used to keep static content on the TLI server to a manageable size. See "Configuring the Scheduling Service" in the *IBM Tealeaf CX Configuration Manual*.

In the Tealeaf Status report, you can review the status of TLI extract jobs in the Extractor Service Log section. See “Tealeaf Status Report” on page 231.

Rolling Archives

New static content detected in the capture stream is written to the current day's `.TLI` file. By default, at 12:01am the daily `.TLI` file is closed, and a new `.TLI` is created for the static objects detected in the coming day.

- The time when the static archive is rolled is configured through the TLI session agent. See "TLI Session Agent" in the *IBM Tealeaf CX Configuration Manual*.
- After a static archive has been rolled, the TLI session agent begins adding content immediately to the new one. It does not scan the daily static archives that have been saved to the backup directory for matching objects.

  Note: When a Canister LSSN file is rolled, the corresponding content in any static archive is not rolled. These rolling operations are managed independently.

Merge operations can be configured to roll the daily static archive content into the monthly static archive. Optionally, the daily static archive that is merged can be moved to the TLI backup folder.

- The TLI backup folder is specified within the TLI session agent. See "TLI Session Agent" in the *IBM Tealeaf CX Configuration Manual*

For more information on configuring tasks for TLI files, see "Configuring the Scheduling Service" in the *IBM Tealeaf CX Configuration Manual*.

Merging Archives

Through the Portal, you can configure daily jobs to merge the contents of each day's static archive into the monthly container archive. This archive can then be preserved for long-term storage.

- During merge operations, static content in the daily archive that is duplicated in the monthly archive is stored in that archive as a reference to prevent duplicated content.
- When a TLI merge operation is configured to be executed each day, the monthly TLI file contains all of the static content for the previous days of the month.

  Note: Monthly archives are rolled at the rolltime on the first day of each month. Monthly archives remain on the TLI server and are the responsibility of the customer to maintain. See "Configuring the Scheduling Service" in the *IBM Tealeaf CX Configuration Manual*.

Deleting Static Content from an Archive

Note: Deletion of individual objects from static archives managed by TLI server is not currently supported.

You may configure trim operations to remove daily static archives from the TLI server that have aged more than a specified number of days through the Scheduling Service. See "Configuring the Scheduling Service" in the *IBM Tealeaf CX Configuration Manual*.

- You may review and delete content from static archives stored locally by RTV. See "Using Static Archives in RTV" in the *IBM Tealeaf RealiTea Viewer User Manual*.

Loading Content from a File

Using an external command-line utility, you can load a static archive file with content referenced through a text file of URLs and optionally the destination static archive.

- If no destination is specified, the static content in the source URLs is loaded into today's archive.

Example Text Files

You may use the URL files to update today's TLI file or a TLI file whose path you specify in the second column. Below, you can review the format for each type of loading.

Note: If you are loading a static archive other than today's file, you must include the full path on the TLI Server to the destination file.

Example URL List for updating a specified TLI file on the TLI server

You can load a specified set of URLs into static archives that are stored on the TLI server. Each line should contain one URL, followed by a `TAB` character and the full path to the local TLI file.

```
http://tealeaf.com/content/PDF/1.pdf      C:\website\content\myPdfSource1.pdf
http://tealeaf.com/content/PDF/2.pdf      C:\website\content\myPdfSource2.pdf
http://tealeaf.com/content/PDF/3.pdf      C:\website\content\myPdfSource3.pdf
http://tealeaf.com/content/PDF/4.pdf      C:\website\content\myPdfSource4.pdf
http://tealeaf.com/content/PDF/5.pdf      C:\website\content\myPdfSource5.pdf
```

Example URL List for updating today's TLI

Each line should contain one URL. No destination file is specified.

```
http://www.example.com/static_assets/SFA_Program.pdf
http://www.example.com/static_assets/lf_map.gif
http://www.example.com/static_assets/IA.png
http://www.example.com/static_assets/IAA.png
http://www.example.com/static_assets/stylesheet.css
```

Executing the Utility

Please complete the following steps to use the utility to execute the load:

1. Login to the TLI Server as an administrator.
2. Copy your text file of URLs to load to the following directory:

   `<Tealeaf_install_directory>\Tools`
3. In the same directory, execute the `TLIUtil.exe` utility using the following command:

   `TLIUtil.exe <file.txt>`

   where:

   - `<file.txt>` is the filename of the text file you stored in the directory.

4. After the load has completed successfully, please delete the text file or move it to a storage location outside of the Tealeaf install directory.
   - To verify the load, you can review the contents of a TLI file through RTV. See "Using Static Archives in RTV" in the *IBM Tealeaf RealiTea Viewer User Manual*.

Replay using Static Archives

This section describes how to use static archives with BBR and RTV.

BBR

After the Replay Server has been configured to use a static archive, sessions delivered through BBR automatically use content stored in appropriately TLI. If Replay Server cannot connect to the TLI Server or cannot locate the appropriate static object, it refers to the origin server to acquire the content.

In the Page Load Details screen, you can identify which objects have been sourced by the Replay Server from the TLI Server. In the Source column:

Value Description

`RemoteTLI`
Object was retrieved from the remote TLI Server by the Replay Server

`TLICache`
Object was retrieved from local TLI cache maintained on the Replay Server

See "CX Browser Based Replay" in the *IBM Tealeaf cxImpact User Manual*.

Note:

Tealeaf administrators can review the static archives currently stored on the TLI Server through the Portal.

- See “Managing Tealeaf Servers” on page 115.

RTV

After RTV is configured to use static archives from the TLI server, no further configuration is required. RTV interacts with Search Server to acquire content through the TLI server. When Search Server delivers the static objects, they are stored in the local IE cache and managed by RTV accordingly.

- See "Using Static Archives in RTV" in the *IBM Tealeaf RealiTea Viewer User Manual*.

Monitoring Static Archives

Tealeaf provides multiple mechanisms for monitoring the content in static archives, the current status of archive creation, and the overall throughput of the system.

- For more information on monitoring static archive content, see “Viewing Static Archive Contents” on page 388.
- For more information on monitoring the static archive pipeline, see “Monitoring Static Archive Pipeline Status” on page 388.
- For more information on monitoring system throughput, see “Static Archives in System Status” on page 388.
- On an ad-hoc basis, you can query all servers in the Tealeaf system, including the TLI Server, to report current status. See “Archive Status in Tealeaf Status” on page 388.

Viewing Static Archive Contents

Through the IBM Tealeaf CX RealiTea Viewer, you can inspect the contents of a TLI archive, including the size, checksum value, URL, and ContentType of each object in the archive.

- This table data can be exported to Excel for further analysis and reporting.
- See "Using Static Archives in RTV" in the *IBM Tealeaf RealiTea Viewer User Manual*.
- IBM Tealeaf CX RealiTea Viewer is a standalone Windows application for searching and replaying sessions. See "RealiTea Viewer (RTV) User Manual" in the *IBM Tealeaf RealiTea Viewer User Manual*.

Monitoring Static Archive Pipeline Status

In the Tealeaf Management System, the Pipeline Status tab displays the session agents in each pipeline and at two-second intervals the current rate of hits and characters processed by each session agent per second.

Steps: To configure the Pipeline Status tab, please complete the following steps.

1. Login to the Tealeaf Portal.
2. From the Portal menu, select **Tealeaf** > **TMS**.
3. The Tealeaf Management System is displayed. Click the Pipeline Status tab.
4. From the Server drop-down, select the server where the TLI pipeline has been created.
5. In the list of pipelines on the right side of the screen, select the TLI pipeline.
6. Processing information for each session agent in the pipeline is displayed in the table on the right side of the screen.
7. To monitor the number of objects that are being inserted into the current `.TLI` file, check the Hit Count column. When this field's value is incremented, another static object has been inserted into the `.TLI` file.
   - See Chapter 6, “TMS Pipeline Status Tab,” on page 297.
   - See Chapter 4, “Tealeaf Management System,” on page 259.
8. If you have configured the TLI processing pipeline to receive hits as directed by an instance of the RTA Split session agent in another pipeline, you may want to review the status of that pipeline, too.

Static Archives in System Status

Through the Tealeaf Portal, you can access multiple reports that can provide insight into operations on static file archives.

Depending on how your TLI pipeline is deployed, the following reports may be of value to you:

- See “System Status” on page 8.
- See “System Status” on page 8.
- See “System Status” on page 8.
- See “System Status” on page 8.

Archive Status in Tealeaf Status

TLI server information can be reviewed through the Tealeaf Portal using a Tealeaf Status report.

- If a TLI Server entry has been added to the Portal Management page, server information is included in the report. See “2. Configure a TLI Server Reference” on page 381.
- You can generate a Tealeaf Status report from the Portal Management page. See “Portal Logs” on page 227.

In the report, the TLI Files section contains information on the TLI files stored on the TLI server registered in the environment.

Field **Description**

Name Filename of the TLI file

Objects
Count of objects in the TLI file

File Size
Size of the TLI file in megabytes

First Use
Timestamp when the first object was inserted into the TLI

Last Use
Timestamp when the most recent object was inserted into the TLI

For more information on status reporting, see “Tealeaf Status Report” on page 231.

Logging

Static archive log messages are generated and stored by the TLI session agent. See "TLI Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

When a TLI extract job has been executed, logging information is available through the Portal. See “Tealeaf Status Report” on page 231.

Managing TLI Files

Over time, the number of TLI files stored on the TLI Server can consume too much space. Through the Scheduling Service, you can configure the trim and merge jobs to manage the total data footprint of static content on the TLI Server.

- See "Configuring the Scheduling Service" in the *IBM Tealeaf CX Configuration Manual*.

Reporting

Static archive report data can be downloaded from RTV and imported into Excel.

- Tealeaf provides an Excel report for summarizing static archive report data. Data exported from RTV can be imported into this preconfigured report for summary and detail reporting on the contents of your static archives.
- See "Using Static Archives in RTV" in the *IBM Tealeaf RealiTea Viewer User Manual*.

Reference

- See “Managing Tealeaf Servers” on page 115.
- See "TLI Session Agent" in the *IBM Tealeaf CX Configuration Manual*.
- See "Using Static Archives in RTV" in the *IBM Tealeaf RealiTea Viewer User Manual*.

Chapter 10. Database Administration

The IBM Tealeaf CX datastore runs on top of a set of four Microsoft SQL Server databases. Tealeaf components access these databases using one of two SQL logins created automatically during installation. These accounts are granted membership to a select set of database roles on each of the Tealeaf databases. They do not have access to any other databases or to the SQL Server instance itself.

- Tealeaf databases can be installed independent of the main Tealeaf installer. See "Installing Tealeaf Databases" in the *IBM Tealeaf Databases Guide*.

Databases Administration Topics

- "Database Backup Strategy" in the *IBM Tealeaf Databases Guide*
- "Database Sizing" in the *IBM Tealeaf Databases Guide*

For more information on troubleshooting database issues, see "Troubleshooting - Tealeaf Databases" in the *IBM Tealeaf Troubleshooting Guide*.

© Copyright IBM Corp. 1999, 2014

Chapter 11. SQL Server authentication

This page contains information about the authentication methods and login permissions that are required by Tealeaf for accessing the SQL Server databases.

Authentication

During the installation process for SQL Server, you are prompted for the type of authentication in use for the install account.

The account in use should have both Windows and SQL Server authentication.

Note: If your account is restricted to Windows-only authentication, contact Tealeaf http://support.tealeaf.com.

Note: The accounts that you use for the Tealeaf databases should not be shared with any other application in your SQL Server environment.

Note: If these logins are created with names other than the default ones, that information must also be updated in the Report Server configuration. See "Configuring the Report Server" in the *IBM Tealeaf CX Configuration Manual*.

Required SQL Server login permissions

The Tealeaf application requires three SQL Server logins to operate correctly:

1. During installation, a SQL Server login is required with sufficient permissions to create and modify the Tealeaf databases. This login is used only during installation. See "Installation login."
2. An Administrative login enables background processes to modify the database schema during normal operation. See "Administrative login" on page 394.
3. A User login is used by all customer-facing portions of the application and only permits read and write access to the databases. See "User login" on page 394.

Installation login

The SQL Server login that is used during installation needs the following permissions to install and modify the schemas of each Tealeaf database.

This login can either be the NT account of the user running the Tealeaf installation or a separate SQL Server login that is provided to the Tealeaf Database Manager as a secondary step after the main installation completes.

Note: The Installation login is required during initial installation of the databases only. After the installation completed successfully, this login can be deleted or disabled.

The simplest option is to use an NT Account that has SQL Server admin privileges or provide a SQL login with that level of access. If that is not possible a login with the following roles and permissions is required to successfully complete the Tealeaf database installation:

- Required Database Roles:
  - `db_ddladmin`

© Copyright IBM Corp. 1999, 2014

  - db_datareader
  - db_datawriter
- Required Database Permissions:
  - ALTER ANY LOGIN
  - ALTER ANY USER
  - CREATE DATABASE
  - CREATE SCHEMA

Administrative login

This login is used by Tealeaf applications that must be able to modify the schemas of one or more of the Tealeaf databases. The applications that use this account that is run in the background and do not allow user interaction.

- Default Account Name - TLADMIN
- Required Database Roles:
  - db_ddladmin
  - db_datareader
  - db_datawriter
- Optional Database Permissions:
  - VIEW SERVER STATE is useful for performance monitoring.

Extra admin login permissions

Through the Portal database administrators can access useful reports on the state of the Tealeaf databases. See "System Status" in the *IBM Tealeaf cxImpact Administration Manual*.

To acquire all of the information to populate the Database Filegroup Size report, the TLADMIN account requires extra permissions. Below are the permissions to enable:

```
USE MASTER
GO

EXEC sp_grantdbaccess 'TLADMIN';

GRANT EXECUTE ON sys.sp_OACreate TO TLADMIN;
GRANT EXECUTE ON sys.sp_OADestroy TO TLADMIN;
GRANT EXECUTE ON sys.sp_OAGetErrorInfo TO TLADMIN;
GRANT EXECUTE ON sys.sp_OAGetProperty TO TLADMIN;
```

If these permissions are not enabled or are not possible, the Database Filegroup Size contains a significant number of zeros in the data.

User login

This login is used by the Tealeaf Portal and any other Tealeaf applications that require database access and allow user interaction. The login is limited to reading and writing data from the Tealeaf databases.

- Default Account Name - TLUSER
- Required Database Roles:
  - db_datareader
  - db_datawriter

ssclr_login login

This login is used by a stored procedure for IBM Tealeaf cxResults to perform enhanced searches against the Visitors database.

- Default Account Name - `ssclr_login`
- Required Database Roles:
  - `db_datareader`
  - `db_datawriter`
  - `EXTERNALACCESS`
  - `ASSEMBLY`

Tealeaf database security

- You can set the default database for the Tealeaf logins `TLADMIN` and `TLUSER` to any valid value.
- No stored procedures in any Tealeaf databases should have access that is granted to public. All access should all be limited explicitly to `TLADMIN` and `TLUSER`. Revoking permissions from public does not cause an issue.
- Many system stored procedures in `TL_SYSTEM` database have run privileges that are granted to public. Revoking these permissions should not cause problems. However, since revoking these permissions requires changes to the MS SQL Server system procedures, Tealeaf cannot guarantee that the changes does not result in system issues. The Tealeaf SQL Server logins are explicitly granted the permissions that they need to access any system-stored procedures.

Reference

SQL Server Administration Topics

- SQL Server Authentication
- "Upgrading SQL Server for Tealeaf" in the *IBM Tealeaf Databases Guide*

Database Administration Topics

- "Database Backup Strategy" in the *IBM Tealeaf Databases Guide*
- "Database Sizing" in the *IBM Tealeaf Databases Guide*

Chapter 12. Configuring the Tealeaf Data Service

The Tealeaf Data Service manages connections between Tealeaf components and services and the databases they are querying. This service allows database connectivity to be managed as a Windows-based service for security, reliability, and control.

- As of Release 8.0, the Reporting Service is an integrated component of the Tealeaf Data Service.

To acquire data, Tealeaf servers and services send queries through the Tealeaf Data Service, which then issues queries to the appropriate databases, canisters, or both.

- Search Server does not communicate directly with the database.
- Any non-Portal Server issuing a command that requires the Tealeaf Data Service is redirected to the Search Server associated with the Portal Server. This instance of Search Server queries the Tealeaf Data Service and returns the data.

Connections

The Tealeaf Data Service replaces some of the functionality managed by the Search Server, which utilizes SQL authentication for direct access to the Tealeaf databases. In some environments, this authentication method is problematic.

Currently, the Tealeaf Data Service manages the following connections:

1. **Tealeaf user data:** Tealeaf user information is made available to Tealeaf components through the Tealeaf data service.
   - See "CX User Administration" in the *IBM Tealeaf cxImpact Administration Manual*.
   - See "cxView User Administration" in the *IBM Tealeaf cxImpact Administration Manual*.
   - See "cxReveal User Administration" in the *IBM Tealeaf cxReveal Administration Manual*.
2. **Tealeaf Status:** The Tealeaf Status report uses Tealeaf Data Service to retrieve statistical information about the Tealeaf servers and databases. See "Tealeaf Status Report" in the *IBM Tealeaf cxImpact Administration Manual*.
3. **Portal Control Settings:** Configuration settings internal to the Portal are stored in the `TL_SYSTEM` database and are retrieved at startup for use.
4. **Tealeaf Canister Server Definitions:** The Portal uses the service to acquire the definitions for all Canister servers in the environment. See "Managing Tealeaf Servers" in the *IBM Tealeaf cxImpact Administration Manual*.
5. **RTV Search & Replay Audit Logs:** The IBM Tealeaf CX RealiTea Viewer uses the service to commit audit logs to the Portal's database which can be seen in the User Activity Reports. See "Monitoring User Activity" in the *IBM Tealeaf cxImpact Administration Manual*.
6. **NT/Active Directory User Lists:** When NT/AD authentication is enabled for Tealeaf, the Portal retrieves the list of configured NT/AD groups from the Data Service. The Data Service collects them at a scheduled interval from Search Servers in the Tealeaf system.

© Copyright IBM Corp. 1999, 2014

7. **Portal Reports:** Wherever possible, the Portal relies on the Tealeaf data service as the data source during report generation. In particular, the data service provides inputs to the following reports:
   - Tealeaf system status reports. See "System Status" in the *IBM Tealeaf cxImpact Administration Manual*.
   - Active status and events reports and the Alert Monitor. See "Monitoring Current Tealeaf Activities" in the *IBM Tealeaf cxImpact User Manual*.
   - The Reporting Service is one component of the Tealeaf Data Service.
8. **Event and Dimension Definitions:** Events, dimensions, session attributes, and other run-time data definitions are managed through Search Server, which queries the Data Service to interact with the database tables.
   - See "Tealeaf Event Manager" in the *IBM Tealeaf Event Manager Manual*.

Installation

During the installation process, the Tealeaf Data Service is installed on the same server where the Tealeaf Data Collector is installed. The Tealeaf Data Service uses the same registry settings to self-configure and connect to the Tealeaf databases.

Database Configuration

For more information on how to configure Tealeaf databases, including re-installation as needed, see "Database Administration" in the *IBM Tealeaf Databases Guide*.

Configuring Tealeaf Data Service

Through Search Server, Tealeaf servers and services communicate with the Data Service and the Report Server by extension.

On each Tealeaf server specified in TMS, the Search Server configuration must be supplied with a method for connecting to the Tealeaf Data Service, depending on the server type:

1. **Portal Server:** The machine hosting the Portal Server must be provided a valid entry for `TeaLeaf Data Service`.
   - In most configurations, the service is installed on the Report Server, which is typically installed on the Portal Server machine. In these configurations, this value is `localhost:23000`.
   - If the Portal Server and the Report Server are split onto separate machines or if the service is installed in a non-standard server, then the value must be specified to point to the hostname and port number of the hosting server:

     ```
     <Server_hosting_Tealeaf_Data_Service>:23000
     ```
2. **All other servers:** All other Tealeaf servers require a valid entry for the Portal Server in their Search Server configurations.
   - These Tealeaf servers communicate through the Portal Server, which knows how to contact the Tealeaf Data Service based on the configuration in the previous step.
   - In Search Server configuration for non-Portal Server machines, the value for `TeaLeaf Data Service` should be left blank.

- For more information on configuring Search Server, see "Configuring the Search Server" in the *IBM Tealeaf CX Configuration Manual*.

- For more information on using TMS, see "TMS WorldView Tab" in the *IBM Tealeaf cxImpact Administration Manual*.

Starting and Stopping the Service

You can start, stop, and restart the Tealeaf Data Service through TMS.

Note: Since the Tealeaf Data Service provides critical data to Search Server, it must be started before Search Server and stopped after it.

1. Log in to the Tealeaf Portal as an administrator.
2. From the Portal menu, select **Tealeaf** > **TMS**.
3. Click the WorldView tab.
4. Select `Servers` view.
5. Click the Data Service node.
6. The available commands are displayed in the Component Actions panel. See "TMS WorldView Tab" in the *IBM Tealeaf cxImpact Administration Manual*.

Logging

This section describes how to manage and use the Data Service log files through IBM Tealeaf Portal.

Accessing Logs

Tealeaf administrators may access the Data Service log files through the Tealeaf Portal.

1. From the Portal menu, select **Tealeaf** > **Portal Management**.
2. Click the Tealeaf Servers link in the left navigation pane.
3. Click the Manage Servers link.
4. The right side of the screen is populated with the configured Tealeaf servers.
5. Select the Report Server.
6. In the toolbar above the server configuration panel, click the Tealeaf Logs icon.
7. From the Filter By drop-down, select `Data Service`.
8. The Data Service logs currently stored on the Report Server are listed.
   - See "Managing Tealeaf Servers" in the *IBM Tealeaf cxImpact Administration Manual*.

You may also access the database entries for the Data Service through the Logs section in the Portal Management page. See "Portal Logs" in the *IBM Tealeaf cxImpact Administration Manual*.

Log File Location

Tealeaf Data Service generates log messages in the following log file:

`<Tealeaf_install_directory>\Logs\TLDataService.log`

The log file is rolled daily at midnight. Log files from previous days include a `YYYYMMDD` indicator in the filename.

Changing the Log Level

By default, the Log Level for the Tealeaf Data Service is set to 1, which means that the log file contains only the following types of status messages:

- Startup
- Shut-down
- Verification of connection to all databases

If needed to resolve specific issues with the Tealeaf Data Service, you can change the logging level to a value between 1 and 9. If Log Level 9 is enabled via TMS, the Data Service log file contains individual request information and additional trace/debug statements to assist in resolving issues.

The Tealeaf Data Service shares the same logging controls as other services managed through the Report Server. To change the logging level for it and other Report Server functions, set the Log Level value on the Database tab in the Report Server configuration in TMS.

- See "Configuring the Report Server" in the *IBM Tealeaf CX Configuration Manual*.

Per-Minute Data Service Stats

In addition to core logging functions, the Data Service writes statistical information on its current state to `.CSV` files in the `Logs` directory. These files provide performance information on the Data Service CPU usage, threads, handles, requests, and more.

Files are in the following location:

```
<Tealeaf_install_directory>\Logs\TLDataService-Stats_YYYYMMDD.log
```

This file is updated every minute and is rolled every day. The file contains the following tab-delimited fields.

Field **Description**

`LogTime`
Timestamp for when the entry was written to the file

`DsUptime`
Current uptime for the Data Service in minutes

`ReqsSinceStart`
Requests since the Data Service was last restarted

`ReqsInLastMinute`
Requests to the Data Service in the last minute

`AvgReqsPerMinute`
Average requests to the Data Service per minute since last restart

`CPU` CPU usage by the Data Service at the time the entry was written

`Handles`
Number of handles in use by the Data Service

`Threads`
Number of threads in use by the Data Service

`WorkingSet`
The current amount of physical memory allocated for the Data Service

`PeakWorkingSet`
The maximum amount of physical memory used by the Data Service since last restart

`PrivateMemory`
The current amount of paged memory allocated by the Data Service

PagedMemory
The current amount of paged memory allocated by the Data Service

PeakPagedMemory
The maximum amount of memory in the virtual memory paging file used by the Data Service since last restart

NonpagedSystemMemory
The current amount of nonpaged system memory allocated for the Data Service

PagedSystemMemory
The current amount of pageable system memory allocated for the Data Service

VirtualMemory
The current amount of the virtual memory allocated for the Data Service

PeakVirtualMemory
The maximum amount of virtual memory used by the Data Service since last restart

AvailWorkerThreads
The number of worker threads currently available to the Data Service

AvailCpThread
The number of available asynchronous I/O threads available to the Data Service

Chapter 13. CX databases

This section provides some overview information on each of the IBM Tealeaf CX databases.

Note: The following filegroup reference information applies to new installations of the databases for Tealeaf Release 8.0 or later. Filegroups for upgraded versions of the databases reflect the earlier version's layout with new filegroups added, based on the new scheme.

System database

The system database is required for the configuration and operation of all Portal and Reporting capabilities. It contains all Tealeaf system information that is created through the configuration and use of Tealeaf.

It contains the following data specific to the customer environment:

- Tealeaf user information
- Dashboards definitions
- Scorecards definitions
- Report definitions
- System log information
- Event definitions
  - Step-based event definitions

Filegroups

Filegroup
Description

PRIMARY
Contains SQL Server database-specific information, including temp space and data management inside the database. Does not contain any Tealeaf data.

CONFIG System configuration information contains all Tealeaf system configuration information, user profiles, and report definitions.

INDEX Contains indexes on all tables in the database.

APP_LOGS
Contains log tables for various components of the Tealeaf application.

Sizing

This database is small, typically between 500 MB and 1 GB. It might experience some growth during normal operation because system logging data is written in to the database.

Reporting database

This database contains statistical information about traffic that is captured from monitored customer applications.

© Copyright IBM Corp. 1999, 2014

This information is collected from the Tealeaf canister(s), aggregated, and stored for use in reporting.

Note: All tables containing indexes are stored in the Reporting database, which enables accurate "point in time" snapshots to be made by performing a backup of this single database.

Filegroups

Filegroup
Description

PRIMARY
Contains SQL Server database-specific information, including temp space and data management inside the database. Does not contain any Tealeaf data.

DATA Contains miscellaneous short-term lookup data

CANDATA
Contains the unaggregated data that is collected from the Processing servers

AGGDATA
Contains hourly aggregated event reporting data

AGGINDEX
Contains indexes for the tables in the AGGDATA filegroup

AGGDATALT
Contains daily aggregated event reporting data

AGGINDEXLT
Contains indexes for the tables in the AGGDATALT filegroup

AGGKEY Contains the dimension report group key information that is used to aggregate dimension data

AGGKEYINDEX
Contains the indexes on the tables in AGGKEY

DIMVALUES
Contains tables that hold all dimension values used in reporting

DIMINDEX
Contains the indexes for the tables in DIMVALUES

EVENTDATA
Contains scorecard and top movers data

EVENTINDEX
Contains the indexes for tables in EVENTDATA

PERFDATA
Contains all hourly Performance reporting data

PERFINDEX
The indexes for the data in PERFDATA

PERFDATALT
Contains all daily Performance reporting data

PERFINDEXLT
The indexes for the data in PERFDATALT

Sizing

This database can be large. Its size is determined by the amount of traffic that flows through the website and the length of time for which you want to retain the reporting data.

Note: For more information about sizing any Tealeaf database, contact Tealeaf Professional Services.

Statistics database

This database contains performance information about the IBM Tealeaf CX system itself. It is used to monitor the systems health and to aid in troubleshooting if an issue arises.

Filegroups

Filegroup
Description

PRIMARY
No Tealeaf data. Used by MSSQL Server for temp space and data management inside the database

DATA Tealeaf system statistics data

INDEX Indexes for the database

LOG The database log

Sizing

This database should remain fairly small. Statistics gathering may be turned off by the user, in which case this database should never grow. If statistics are being gathered, then the size of the database is determined by the length of time the data is kept before being trimmed.

Reports

Tealeaf captures statistical information from these databases, which is available through the Portal.

- See "System Status" in the *IBM Tealeaf cxImpact Administration Manual*.
- See "System Status" in the *IBM Tealeaf cxImpact Administration Manual*.

Chapter 14. cxResults databases

Note: The following filegroup reference information applies to new installations of the databases for Tealeaf Release 8.0 or later. Filegroups for upgraded versions of the databases reflect the earlier version's layout with new filegroups added, based on the new scheme.

Session Segment database

This database contains analytical data that is parsed from sessions that are identified by Tealeaf users through the Portal Search and session segment analysis functionality.

The data is used to provide in-depth analysis of groups of sessions pertaining to a particular issue or event.

Filegroups

Filegroup
Description

PRIMARY
Used for short-term storage and for the storage of small lookup tables.

DATA Stores the analytic data that is used for session segment analysis.

INDEX Non-clustered index storage.

Sizing

This database can be fairly large. Its size is dependent on the number and size of session segments being analyzed. If session segment analysis is not being used, this database should experience no growth.

Visitor Staging database

The Visitor Staging database contains the raw data that is extracted from sessions in the Tealeaf canisters.

The data is extracted into this database to provide flexibility in other ETL operations and to separate those operations from the reporting capabilities of IBM Tealeaf cxResults.

Filegroups

Filegroup
Description

PRIMARY
Contains basic database configuration and SQL Server information.

VISITOR
Contains the unique visitor information that is extracted from the canister sessions.

VISITORINDEX
Contains index information for the tables in the Visitor filegroup.

© Copyright IBM Corp. 1999, 2014

SESSION
Contains the raw session information that is extracted from the Tealeaf canisters.

SESSIONINDEX
Contains index information about the tables in the Session filegroup.

EVENT Contains the raw event data that is extracted from the Tealeaf canister sessions.

EVENTINDEX
Contains index information about the tables in the Event filegroup.

DATA Look up and Index data for the VDB extract process

Sizing

Under normal operation the Staging database contains only one week of information. Assuming that traffic levels on the website remain relatively constant, this database should reach a stable size and remain relatively unchanged. The actual size of the database can be fairly large and varies depending several factors. Refer to the database sizing section for further details.

Visitor Reporting database

The Visitor Reporting database contains a snapshot of the visitor data that is contained in the Tealeaf canisters.

This data supports the search and reporting capabilities IBM Tealeaf cxResults application. The data is loaded through an ETL process from the Staging database.

Filegroups

Filegroup
Description

PRIMARY
Contains basic database configuration and SQL Server information.

DATA Contains the segment information that is generated through searches in the IBM Tealeaf cxResults application.

LOOKUP Contains largely static lookup information for use in reporting (for example, event definitions, reference information, and so on).

VISITOR
Contains all unique visitor information.

SESSION
Contains all the extracted session information.

EVENT Contains all the information about events the fired in the customers system.

INDEX Contains all the indexes for the database.

Sizing

The reporting database contains a snapshot of the visitor information that passes through the customers website. The time interval of the snapshot is configurable and is typically 2 - 3 months. As a result, this database can be quite large, but its size should stabilize after the first time interval. At that point, the snapshot

functions like a rolling window, removing older data and replacing it with current data. For exact sizing information, refer to "Database Sizing" in the *IBM Tealeaf Databases Guide*.

Reports

Tealeaf captures statistical information from these databases, which is available through the Portal.

- See "System Status" in the *IBM Tealeaf cxImpact Administration Manual*.
- See "System Status" in the *IBM Tealeaf cxImpact Administration Manual*.

Chapter 15. Database backup strategy

Regularly, your Tealeaf databases should be backed up for recovery and performance reasons. Without valid backups, Tealeaf may be unable to recover customer data in the event of a system failure or data corruption. Without a proper backup strategy, the log files for the Tealeaf databases may grow unbounded, causing both disk space loss and performance issues.

The strategies that are outlined below are recommended for providing the most reliable recovery capability while minimizing storage overhead and potential interference with the operation of the Tealeaf system.

Note: Tealeaf strongly recommends performing regular backups of the Tealeaf databases using the recommended backup strategies in this section. Failure to back up your databases may result in unrecoverable loss of data.

Backup mechanisms

You can back up your Tealeaf databases using either of the following ways:

- The Tealeaf installation process includes pre-scripted SQL files to create and schedule back up jobs consistent with the backup strategy recommended by Tealeaf. See “Database backup scripts” on page 412.
- You can use the Tealeaf Database Manager to configure and run regularly scheduled backups. See "Using Tealeaf Database Manager" in the *IBM Tealeaf Databases Guide*.

Recommended backup strategy

Note: Tealeaf recommends running a nightly Full Backup on all Tealeaf databases.

The recommended backup strategy should be deployed for all Tealeaf databases, whenever possible. These databases contain information essential to the system operation or information that is not easily recovered from another part of the Tealeaf system.

By default, the Tealeaf databases are installed and maintained in Simple Recovery Mode. As a result, incremental backups are not necessary.

Note: Tealeaf recommends the SIMPLE recovery model for SQL Server. If your instance of SQL Server is configured for SIMPLE recovery model, it is not possible to back up the database logs.

Alternate backup strategies

Note: If the Tealeaf databases cannot be backed up on a nightly basis or another recovery model or backup strategy is required for business reasons, you should ensure that the Tealeaf databases are backed up and maintained properly. Failure to do so may result in wasted space and decreased system performance.

© Copyright IBM Corp. 1999, 2014

Database backup scripts

As part of the installation process, a set of database backup scripts is generated according to the strategies listed above.

The scripts are stored in the following directory:

`<Tealeaf_install_directory>\Reporting\DatabaseSchema`

Generated backup script files are of the following pattern: `*DBBackup.sql`

Backup notifications

By default, the backup jobs that are created by the scripts insert a message in the Event Log if the backup job fails to run.

- The reason for the error can be seen by viewing the history of the backup job in SQL Server Management Studio.
- The backup jobs can be configured to send emails on success and failure if the SQL Server Notification Service is installed and enabled. See "SQL Server notifications" on page 413.

Setting up database backups using the scripts

1. Each backup script contains a section of variables at the top of the file. The values of these variables are the only things that should be edited in the file, as follows:

```
/*****************************************************
    * Script Variables
    * @jobName         - The name of the backup job
    * @jobOwner        - The SQL Server Login that will own the backup job
    * @operator        - The name of the email operator who will be notified
    *                    if the job fails (added through sp_add_operator).
    * @scheduleName    - The name of the backup job schedule
    * @database        - The name of the CX   database
    * @fullBackupFile   - The filename of the full backup of the   database.
    *                     By default this will be
    *                     overwritten with a new backup twice a week (Sunday,
    *                     Wednesday).
    * @logBackupFile    - The filename of the log bakup for the   database.
    *                     This will be appended to
    *                    daily and started fresh on Monday and Thursday by
    *                    default.
    *****************************************************/
   DECLARE @database nvarchar(256),
           @fullBackupFile nvarchar(256),
           @jobOwner nvarchar(128),
           @jobName nvarchar(128),
           @scheduleName nvarchar(128),
           @serverName nvarchar(30),
           @operator nvarchar(128),
           @backupCmd nvarchar(4000);

   SELECT  @jobName = N'cxImpact Reporting Database Backup',
           @scheduleName = N'cxImpact Reporting Backup Job',
           @jobOwner = N'sa',
           @database = N'TL_REPORTS',
           @fullBackupFile = N'c:\Tealeaf\db_backups\tl_reports_full.bak',
           @serverName = N'(local)',
           @operator = NULL;

   /********* DO NOT CHANGE ANYTHING BELOW THIS LINE *********/
```

2. The only required change is to assign valid file paths to the `@fullBackupFile` and `@logBackupFile` variables.
3. To create and schedule a backup for a particular database:
   a. Open the file in SQL Server Management Studio. The scripts must be run as SA or an equivalent login.
   b. Set the files paths correctly.
   c. If the SQL Server Notification service is being used, set the `@operator` variable to a valid SSNS Operator. When run, the backup job sends email to the specified address.
      - For more information about configuring the operator, see “SQL Server notifications.”
   d. Run the script.

SQL Server notifications

If the SQL Server Notification Service is installed and enabled, the backup jobs can be configured to send email notifications upon success, failure, or both.

Using SQL Notification Service requires the creation of an Operator. To assist, an extra script, `AddNotificationOperator.sql`, is created during the installation process in the same directory as the backup scripts.

To create the operator:

1. Open the `AddNotificationOperator.sql` script in SQL Management Studio.
2. Set valid display name and email address values.
3. Run the script.

Reference

SQL Server administration topics

- "SQL Server Authentication" in the *IBM Tealeaf Databases Guide*
- "Upgrading SQL Server for Tealeaf" in the *IBM Tealeaf Databases Guide*

Database administration topics

- Database Backup Strategy
- "Database Sizing" in the *IBM Tealeaf Databases Guide*

Chapter 16. Database sizing

Sizing guidelines for Tealeaf databases are based on multiple interdependent factors.
For detailed guidelines, please contact Tealeaf Professional Services.

- For general sizing guidelines for IBM Tealeaf CX databases, see "CX Databases" in the *IBM Tealeaf Databases Guide*.
- For general sizing guidelines for IBM Tealeaf cxResults databases, see "cxResults Databases" in the *IBM Tealeaf Databases Guide*.

Setting the recovery model and logging levels for Tealeaf databases

To properly maintain the Tealeaf databases, control their size, and improve performance, Tealeaf recommends using the Simple Recovery Model for SQL Server and compressing the log files.

If you are experiencing database size and performance issues, review the following steps to set the Recovery model to the `Simple` setting:

1. For each installed Tealeaf database:
   a. In SQL Server Management Studio, right-click each Tealeaf database and select **Properties**.
   b. In the dialog, select the **Options** page. Review the setting in the **Recovery Model** drop-down. To reduce the size of the database, set this setting to `Simple`. Click **OK** at the bottom of the dialog.
   c. Repeat the above steps for the other Tealeaf databases.
2. Then, stop the Tealeaf Data Collector and Tealeaf Data Service through the Windows Services Control Panel.
3. In SQL Server Management Studio, for each database:
   a. Right-click the database name. Select **Tasks** > **Shrink** > **Files**.
   b. In the dialog, select `Log` from the **File Type** drop-down. Click **OK**.
   c. Repeat the above steps for the other Tealeaf databases.
4. After the above was completed for each database, restart the Tealeaf services.

These steps should reduce the size of your databases and reduce the disk I/O by lowering the amount of logging that is being done.

Reference

SQL Server administration topics

- "SQL Server Authentication" in the *IBM Tealeaf Databases Guide*
- "Upgrading SQL Server for Tealeaf" in the *IBM Tealeaf Databases Guide*

Database administration topics

- "Database Backup Strategy" in the *IBM Tealeaf Databases Guide*
- Database Sizing

© Copyright IBM Corp. 1999, 2014

Chapter 17. Troubleshooting Tealeaf databases

If you are experiencing an access or performance-related issue with your installed Tealeaf SQL database, review the following troubleshooting tips, which may help to resolve your issue.

Note: For some functions of the Tealeaf Database Manager, you must have System Administrator privileges for the SQL Server hosting the Tealeaf databases.

The Tealeaf SQL databases are installed using the Tealeaf Database Manager, which can be used to upgrade or reinstall the databases. Some of these steps may require reinstallation or upgrading. See "Using Tealeaf Database Manager" in the *IBM Tealeaf Databases Guide*.

Note: The Tealeaf Database Manager is used to install the SQL Reporting databases and IBM Tealeaf cxResults databases. It cannot be used to install or update the Canister databases, which are installed as part of the Processing Server through `Setup.exe` in the initial installation. For more information about those databases, see "CX Installation and Setup" in the *IBM Tealeaf CX Installation Manual*.

- For more information about configuring the Canister databases, see "Configuring the CX Canister" in the *IBM Tealeaf CX Configuration Manual*.

SQL Server configuration

Unable to connect to SQL Server

If Tealeaf is unable to connect to the SQL Server database, try the following steps:

1. **SQL Server custom ports**
   - By default, SQL Server listens on port 1433, which is also the default communication port for Tealeaf, unless selected differently during installation. You can verify this value through the SQL Server Configuration Manager.
2. **Unable to connect through dynamic ports**
   - If your SQL Server implementation uses dynamic ports, you must sync your Tealeaf solution with them. See “Configuring use of SQL Server dynamic ports” on page 418.
3. **SQL network protocols**
   - By default, SQL Server does not enable TCP and Named Pipes, which are used by Tealeaf. Verify that these protocols are enabled through the SQL Configuration Manager.

Failure to install due to model database size mismatch

During installation of the Tealeaf databases, the Tealeaf Database Manager may report a failure because of a mismatch between the configured database sizes in Tealeaf and the model database sizes that are configured in SQL Server Management Studio.

In rare cases, a customer may change the size of the model database size in their SQL Server installation from the default value of 50MB. SQL Server Management Studio does not allow databases to be installed below the model database size. For

© Copyright IBM Corp. 1999, 2014

example, if the model database size is set to 250MB, when the Tealeaf Database Manager attempts to install the Tealeaf databases, some installations fail because they are smaller than this minimum limit.

The solution is to do one of the following:

1. Reconfigure the model database size in SQL Server Management Studio to 50MB. For more information, consult the documentation that was provided with your SQL Server product.
2. Reconfigure the database sizes in Tealeaf Database Manager to be slightly larger than the model database size in SQL Server Management Studio. For the above example, you could set the database size to 260MB.
   - Database sizes must be configured through individual filegroups for each database. See "Tealeaf Database Manager Reference" in the *IBM Tealeaf Databases Guide*.

Configuring use of SQL Server dynamic ports

Tealeaf does not natively support use of dynamic ports when accessing SQL Server. If possible, Tealeaf recommends disabling use of dynamic ports.

When SQL Server starts, a dynamic port is selected. This port is used during operations. During restart, the configured port remains, unless a conflict emerges.

If the ports must remain dynamic, complete the following steps.

1. Start all Tealeaf services.
2. Start the database.
3. Through SQL Server Management Studio, locate the port that SQL Server is using.
4. This port must be populated for each database in the Connection dialog of the Tealeaf Database Manager. See "Tealeaf Database Manager Reference" in the *IBM Tealeaf Databases Guide*.
5. In the Tealeaf Database Manager, you can set the ports for the databases.
   a. From the TDM menu, select **Mode** > **Info/Config** > **Report Server Configuration**.
   b. Populate the ports accordingly. When changes are applied, the registry keys below are updated for you.
      - See "Tealeaf Database Manager Reference" in the *IBM Tealeaf Databases Guide*.
6. Through Tealeaf Database Manager, install or upgrade the databases as necessary.
   - See "Using Tealeaf Database Manager" in the *IBM Tealeaf Databases Guide*.

On the Report Server, database port numbers are in the following registry keys. You can update these through the Tealeaf Database Manager Connecting screen.

Product (Database)
Registry Key

IBM Tealeaf cxImpact (Reports)
`Tealeaf Technology > DataStore > Report Server > Port`

IBM Tealeaf cxResults (Visitors)
`Tealeaf Technology > DataStore > Report Server > Visitor Port`

IBM Tealeaf cxReveal (Search)

```
Tealeaf Technology > DataStore > Database > Search > Database Port
```

Note: If SQL Server is forced to choose a new dynamic port, then the new port must be used whenever you access the databases through the Connecting screen in the Tealeaf Database Manager and repopulated in the appropriate registry key through the Tealeaf Database Manager using the above steps.

Migrating to a new SQL Server

If you are migrating to a new instance of SQL Server, there are more steps that are required.

See "Migrating Tealeaf Databases" in the *IBM Tealeaf Databases Guide*.

Restore failed when MDF file name is claimed

When restoring databases to a new instance of SQL Server, you may encounter an error similar to the following:

```
Restore failed for Server <MyServer>

Additional information:System.Data.SqlClient.SqlError: File:
'C:\Program Files\Microsoft SQL Server\
MSSQL10_50.MSSQLSERVER\MSSQL\DATA\RL_REPORTS.mdf' is claimed by 'DATA'(3) and
'PRIMARY'(1). The WITH MOVE clause can be used to relocate one or more files.
(Microsoft.SqlServer.Smo)
```

The above issue is caused by the SQL Server restore procedure that converts all of the Tealeaf database names to `RL_REPORT.mdf`. This situation is typically caused when the new location has a different directory structure or when the databases are moved to a new version of SQL Server.

The solution is to manually append the file name of each data file to include the Filegroup identifier for the file to the end of the file name.

For more information about the appropriate steps to restore the databases, see "Migrating Tealeaf Databases" in the *IBM Tealeaf Databases Guide*.

Poor SQL Server performance on 64-Bit operating systems

Note: If you are using Tealeaf version 8.8 or later, SQL Server 2005 is no longer supported. For more information, see "Supported SQL Server Versions" in the *IBM Tealeaf CX Databases Guide*.

Tealeaf databases are designed to take advantage of 64-bit operating system capabilities. However, because of a Microsoft issue, performance issues may arrive when running SQL Server 2005 on a 64-bit version of Windows. These issues can include the following:

- The performance of SQL Server 2005 decreases suddenly.
- SQL Server 2005 stops responding for a short time.
- A timeout occurs for applications that connect to SQL Server 2005.
- Problems occur when you run even simple commands or use applications on the system.

The issue may be that 64-bit Windows operating system is paging out the working set of the SQL Server process. When the SQL Server process reaches 50 percent of

the memory that is allocated for the process on SQL Server 2005 SP2 or later, error messages may begin appearing in the SQL Server error log.

To fix the issue:

1. Review the Microsoft troubleshooting steps. For more information, review "How to troubleshoot this problem" on `http://support.microsoft.com/kb/918483`.
2. If the problem persists, you can prevent Windows from paging out the buffer pool memory of the SQL Server process. You can lock the memory by assigning the `Lock pages in memory` user right to the user account that is the startup account for the SQL Server service.
   - This step prevents the operating system from paging out memory and instead allows SQL Server to do the paging. For more information, see `http://technet.microsoft.com/en-us/library/ms190730.aspx`.
3. Restart the SQL Server.

For more information about this topic including the troubleshooting steps, visit `http://support.microsoft.com/kb/918483`.

SuperSocket information: (SpnRegister): error 1355 in Windows application event log

From http://support.microsoft.com/kb/303411:

> This message is not an error message. This text is only a warning that SQL Server was not able to register a Service Principal Name (SPN), which indicates that the security mechanism used is Microsoft Windows NT Challenge\Response (NTLM) authentication instead of Kerberos.
> CAUSE
> The message usually appears because the SQL Server service account is running as a domain user who does not have requisite permissions to register SPNs.

SQL Server blocks access to procedure sys.sp_OAGetProperty

In some environments, the following error may be displayed in the SQL Server error log:

```
SQL Server blocked access to procedure 'sys.sp_OAGetProperty' of component
'Ole Automation Procedures' because this component is turned off as part of
the security configuration for this server. A system administrator can
enable the use of 'Ole Automation Procedures' by using sp_configure.
For more information about enabling 'Ole Automation Procedures',
see "Surface Area Configuration" in SQL Server Books Online.
```

`sys.sp_OAGetProperty` is an automation procedure that allows access to system details. Tealeaf uses this stored procedure to acquire details on disk space.

Note: If access to the `sys.sp_OAGetProperty` procedure is not enabled:

1. The Database Filegroup Size and Database Table Size reports in the Portal do not contain free space and unused space information.
2. Available disk space on SQL Server is not recorded and does not appear in the Tealeaf Event Log.

The error message is harmless. However, to avoid confusion, you may want to deny access to pr_ServerDiskSpace, which uses this stored procedure. When access is denied, error messages are reported from pr_ServerDiskSpace, instead of sys.sp_OAGetProperty. Error messages from the latter procedure may suggest possible hacks or malicious software, so unnecessary messages should be suppressed.

To deny access, run the following through SQL Server Management Studio:

```
use TL_STATISTICS
go
Deny execute on pr_ServerDiskSpace to TLAdmin;
Deny execute on pr_ServerDiskSpace to TLUser;

use TL_visreport
go
Deny execute on pr_ServerDiskSpace to TLAdmin;
Deny execute on pr_ServerDiskSpace to TLUser
```

Tealeaf database configuration

Unable to connect

If you are unable to access the Tealeaf database, you should verify that you are using the appropriate full-qualified host name.

If you are running against a named database instance, then the host name must follow a specific format:
<host_name>\<instance_name>
where
<host_name> is the IP address, the machine name, or localhost
<instance_name> is the name of the DB instance.

See "Using Tealeaf Database Manager" in the *IBM Tealeaf Databases Guide*.

During upgrade, Tealeaf Database Manager fails to upgrade the Result Set Extractor

The required views do not exist. The RSE Database must be re-created. See “Result Set Extractor database fails to analyze a result set” on page 423.

Failing to create database in clustered server environment

During installation into a clustered server environment, if you are seeing error messages during creation, you may have improperly specified your SQL path to a disk that is not part of the cluster group or dependency list.

To resolve the cluster dependency issue, see http://support.microsoft.com/kb/295732.

In the Tealeaf Database Manager, you should also verify the path where the database files are stored. If you do not know the path to a remote server, contact the appropriate person.

Note: The path must be the absolute path relative to the SQL Server. Network paths are not permitted.

See "Using Tealeaf Database Manager" in the *IBM Tealeaf Databases Guide*.

Tealeaf database performance

Memory allocation

If any of the following conditions are occurring:

1. Tealeaf is performing slowly
2. Reports are delayed in arriving
3. The box running SQL Server is having paging or disk I/O problems

The memory that is allocated for SQL Server may be set too low. You can verify and change the amount of memory that is allocated to the Tealeaf databases in the SQL Server Management Studio.

Fragmentation

Normal database operations may cause the hard disk drive media and the stored Tealeaf databases to become fragmented. Over time, disk fragmentation can significantly affect performance.

- As part of regular server maintenance, your DBA should perform disk fragmentation checks and defragment disks periodically.

The Tealeaf database indexes can also become fragmented. Periodically, these should be defragmented to ensure optimal performance.

Note: Generally, database index fragmentation should be kept below 20%.

- **Release 7.1 or earlier:** You can run fragmentation checks and defragmentation tasks through Microsoft SQL Server Administration Console. For more information, consult the documentation that came with the product.
- **Release 7.2:** Through the Tealeaf Database Manager, you can check the fragmentation status of your database indexes and run defragmentation tasks as needed. See "Using Tealeaf Database Manager" in the *IBM Tealeaf Databases Guide*.

Tealeaf Database Manager hangs when upgrading the System database

The Tealeaf Database Manager may hang when the System database is being upgrade if other components of the system are running.

While this situation may be caused by multiple factors, in many cases the cause is the presence of an active Stats Logger session agent in the Windows pipeline of the Report Server. This session agent is used to collect statistical information in the pipeline and submit it to the Statistics database for use in Portal-based reports.

- See "Statistics Logger Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

The active session agent may be keeping the Statistics database open and therefore unavailable to the Tealeaf Database Manager, which causes it to hang.

Possible solutions

To resolve this issue, you may pursue any of the following solutions:

1. Stop the Transport Service before you start the Database Manager.

   Note: As soon as you stop the Transport Service, hits are dropped and data is lost.

   a. After you use the Database Manager, you must remember to restart the Transport Service through TMS.

- See "Configuring the Transport Service" in the *IBM Tealeaf CX Configuration Manual*.
- See "TMS WorldView Tab" in the *IBM Tealeaf cxImpact Administration Manual*.

2. Remove or disable the Stats Logger session agent from your processing pipeline.

   Note: Tealeaf recommends disabling the Stats Logger session agent during Database Manager operations in an All-in-One system.

   a. When the session agent is removed from the pipeline, you may start the Database Manager.
   b. Remember to add it back through TMS after you exited the Database Manager.

   - See "TMS Pipeline Editor" in the *IBM Tealeaf cxImpact Administration Manual*.

3. If the Stats Logger session agent is part of a dedicated child pipeline in the Canister, you can disable the whole pipeline through TMS.
   - See "TMS Pipeline Editor" in the *IBM Tealeaf cxImpact Administration Manual*.

Issues with individual databases

Visitor database for cxResults

Note: The Visitor database applies to IBM Tealeaf cxResults. IBM Tealeaf cxResults is no longer available as a newly licensed product as of Release 8.7. Customers that licensed IBM Tealeaf cxResults in Release 8.6 and earlier may continue to use and receive support for the product in Release 8.7 and later. For more information, contact Tealeaf Customer Support .

Data connection timeout in Visitor database

Although there is a data connection timeout setting for the Visitor database, the default value is set high. It should not need to be changed.

Result Set Extractor Database for cxResults

Result Set Extractor database fails to analyze a result set

During an upgrade or reinstall of the RSE database, you may encounter the following error message, which is followed by a table create statement:

```
Failed to create the database views:
Invalid object name 'RSE_HITSTATISTICS'.
```

In this case, the RSE database may be missing some views. In this case, the RSE database empty. To resolve this issue, the database must be re-created:

1. Run the Tealeaf Database Manager. See "Using Tealeaf Database Manager" in the *IBM Tealeaf Databases Guide*.
2. Select **Advanced mode**.
3. In the Database Install tab, click **Create**.
4. In the Database Type pane, click **RSE**.
5. Click **Run**. The RSE database is re-created.
   - Review and observe any generated error messages.
6. After the reinstallation is complete, you may upgrade the database. See "Using Tealeaf Database Manager" in the *IBM Tealeaf Databases Guide*.

Search database for cxReveal

cxReveal database install forces collation setting

If you attempted to install the IBM Tealeaf cxReveal search database, the database collation setting may be forced to the following:

```
SQL_Latin1_General_CP1_CI_AS
```

This setting is configured regardless of the collation setting in SQL Server, which presents a problem for international customers, specifically those in Europe.

The solution to drop and reinstall the TL_SEARCH database through the Tealeaf Database Manager.

Note: Before you begin, verify that the System database (TL_SYSTEM) and the Reports database (TL_REPORTS) were installed or upgraded to the version for which you are installing the TL_SEARCH database.

1. On the Portal Server, start the Tealeaf Database Manager.
2. From the TDM menu, select **Mode** > **Database Setup** > **Single Database Mode**.
3. Drop the Search (TL_SEARCH) database:
   a. From the Single Database Options screen, select Search from the Database drop-down.
   b. From the Action drop-down, select Uninstall.
   c. Click **OK**.
   d. The database is uninstalled.
4. Create the Search database:
   a. Use the Search database creation script, which is provided in the following location:
      <Tealeaf_install_directory>\SQL\DBCreateScripts\CREATE_TL_SEARCH.sql
   b. Start the Tealeaf Database Manager.
   c. From the TDM menu, select **Mode** > **Database Setup** > **Single Database Mode**.
   d. From the TDM menu, select **DB Creation** > **Create Databases** so that the option is not selected.
   e. From the Single Database Options screen, select Search from the Database drop-down.
   f. From the Action drop-down, select Install.
   g. Click **OK**.
5. The Search database is installed with the proper collation setting.
6. If you are performing further installation operations, you might want to re-enable the Create Databases option. From the TDM menu, select **DB Creation** > **Create Databases** so that the option is selected.

Configuration issues

Data collection

Poor data collection performance

If the data collection process is consuming large volumes of memory or failing to finish, you may need to adjust the Data Collector Batch Size setting, which defines

the maximum number of records to extract or load for the Data Collector in a single batch. The default value is set to `2000`.

Note: Do not set this value over `5000`.

See "CX Settings" in the *IBM Tealeaf cxImpact Administration Manual*.

Slow data collection times

Typically, data collection should run in under 5 minutes. If the data collection process is more than 10 minutes, then you should increase the size of the Data Trim Interval setting.

The default value is `1 - Hourly`. Setting it to a larger value causes this process to run less frequently. However, you should see temporary spikes in the size of your database.
See "CX Settings" in the *IBM Tealeaf cxImpact Administration Manual*.

Database connection timeouts

If your Tealeaf databases are experiencing repeated timeouts, you may need to adjust the Database Connection Timeout setting. The default connection timeout is set to `30` seconds. After the system was running for a while and traffic was increased, your hardware may not be able to manage the traffic level at the current setting.

Try setting this value to double its current value. If it happens again, double it again.

Note: If you must set the connection timeout above 300 seconds (5 minutes), you may have more issues in your environment.

See "CX Settings" in the *IBM Tealeaf cxImpact Administration Manual*.

Data collection performance from multiple canisters

By default, Tealeaf can collect from two canisters at one time. If your environment has more canisters, you can increase the size of the `Data Collection - Max Concurrent` setting. The appropriate setting depends on your hardware environment.

- See "CX Settings" in the *IBM Tealeaf cxImpact Administration Manual*.

If you notice any of the following:

1. Some slow-down in data collection
2. Errors in the data collection process
3. The canisters are busy

Then you should try decreasing this setting.

Database is getting too large

If the size of your database is pushing the limits of your storage, you may consider adjusting some settings.

1. If you do not use the Page Performance reports, you do not need to collect the Path Statistics. You can disable the Path Statistics collection to save significant database space.
   - See "CX Settings" in the *IBM Tealeaf cxImpact Administration Manual*.

2. In SQL Server Management Studio, check the Recovery Model setting for each database. If the Recovery Model is set to `Full`, the databases and their log files can grow larger. See "Database Sizing" in the *IBM Tealeaf Databases Guide*.

Error - an existing connection was forcibly closed by the remote host

From time to time, the Data Collector may register the following log message:

```
An existing connection was forcibly closed by the remote host
```

This error indicates that SQL Server or the server hosting it closed the connection unexpectedly. Typically, SQL Server was restarted, or an operation that required closing all existing connections was run, such as a backup or restore operation.

When these errors occur sporadically, they do not affect data collection or the validity of the data.

Note: If these errors are occurring frequently, verify that data collection is actually completing. You can review the log information for the most recent Data Collector that is run through the Portal. See "Portal Logs" in the *IBM Tealeaf cxImpact Administration Manual*.

Visitor Database Extractor

The Visitor Database Extractor is used to extract session data from the database and to insert it into the staging tables for the Visitor database. This section provides some suggestions for how to troubleshoot issues with the VDB Extractor.

- See "Configuring the Scheduling Service" in the *IBM Tealeaf CX Configuration Manual*.

Unable to search for visitors

If you are unable to complete searches of the Visitor database or are no longer getting updated results, the Visitor database extractor job may be disabled. For more information about enabling this job, see "Configuring the Scheduling Service" in the *IBM Tealeaf CX Configuration Manual*.

Troubleshooting database upgrades

Invalid filegroup specified error during database upgrade

During the upgrade of one or more Tealeaf databases, you may receive an error similar to the following:

```
Invalid filegroup <name> specified.
```

This error indicates that the database file (`<name>`) is missing.

To fix this issue, you must determine if the database containing the file is supposed to contain a single filegroup or not.

If the database contains a single filegroup

Note: The only Tealeaf database containing a single filegroup is the `TL_REPORTS` database if it was upgraded from a Release 4.x version of Tealeaf. All other Tealeaf databases contain multiple filegroups.

Try reinstalling the database through the Tealeaf Database Manager. See "Installing Tealeaf Databases" in the *IBM Tealeaf Databases Guide*.

If the database contains multiple filegroups, including the missing one:
If the database is not a single filegroup database, a missing filegroup indicates database corruption. This serious issue may require rebuilding the database or restoring from backup.

- For more information, contact Tealeaf http://support.tealeaf.com.

Checklist

If you are still experiencing difficulties with your Tealeaf databases, acquire the following information from your system:

1. Acquire the full specs of the box running the Reporting/SQL Server, including CPU, RAM, and information about each hard disk drive. Drive information should include number of drives, contents of each drive, and available disk space.
2. Acquire the version of Tealeaf that is in use. Improvements may be made in subsequent Tealeaf releases.
3. What is the traffic load? How many pages per day? How many events per page? How many pages per session?
4. Is the entire Portal sluggish or certain items in the Portal?
   - If the latter, the problem is likely caused by issues with the code or configuration of the individual item(s).
5. Set `Log Level = 9`. Acquire a day or two of `TealeafReportingSvc.log` messages.
   - This information is useful for troubleshooting.
6. On the SQL Server, open **Performance Monitor**. Watch the Avg. Disk Queue and Page/Sec counters for 5 to 10 minutes. Are
   either of them consistently pegged through the roof?
   - Does **Performance Monitor** show much paging on the box? This may indicate that the Data Service is starved for memory.
7. How large is the TL_REPORTS database?

When you acquired the above information, open a support ticket with http://support.tealeaf.com.

Chapter 18. Report Access Permissions

Users whose accounts contain the proper permissions may administer Tealeaf reports. Depending on the type of report, permissions to administer individual reports of that type are configured in the report definition.

This section provides information on how configure permissions to administer each type of report available in Tealeaf.

Scorecard Access Permissions

Available through IBM Tealeaf cxView, scorecards enable the graphical display of KPI or process data through the Tealeaf Portal. Scorecards can be scheduled for delivery to selected users.

Note: Scorecards are a component of IBM Tealeaf cxView, a separately licensable component of the IBM Tealeaf CX system. please contact your IBM Tealeaf representative.

- To access the available scorecards, select **Analyze** > **Scorecards** in the Tealeaf Portal.
- For more information on using scorecards, see "Using Scorecards" in the *IBM Tealeaf cxView User Manual*.

General Admin Permissions

Scorecards can be administered by any user with permissions to configure scorecards.

Note: Any individual who can configure scorecards can administer any available scorecard, including deleting it.

- To configure scorecards, select **Configure** > **Scorecards** in the Portal menu.
- If you cannot access this menu item, the menu item must be added to your menu profile. See “cxView User Administration” on page 184 in the “cxView User Administration” on page 184 page.

Access Permissions for Individual Scorecards

Access to individual scorecards is defined in the Configuring Scorecards page.

1. To configure scorecards, select **Configure** > **Scorecards** in the Portal menu.
2. In the left panel, select the type of scorecard you are administering: **Process Scorecards** or **KPI Scorecards**.
3. The list of available scorecards for the selected type is displayed.
4. To administer a scorecard, select it and click the Pencil ( ) icon.
5. Click the Permissions tab.
6. In the Permissions tab, you can configure the user groups that are permitted to see the scorecard.
7. To save changes, click **Save**.
   - See "Configuring Process Scorecards" in the *IBM Tealeaf cxView User Manual*.
   - See "Configuring KPI Scorecards" in the *IBM Tealeaf cxView User Manual*.

© Copyright IBM Corp. 1999, 2014

Dashboard Access Permissions

Available through IBM Tealeaf cxView, dashboards enable the graphical display of scorecard and other reporting components in one or more tabs through the Tealeaf Portal. Dashboards can be scheduled for delivery to selected users.

Note: Dashboards are a component of IBM Tealeaf cxView, a separately licensable component of the IBM Tealeaf CX system. please contact your IBM Tealeaf representative.

- To access the available dashboards, select an item from the Dashboards menu.
- For more information on using dashboards, see "Using Dashboards" in the *IBM Tealeaf cxView User Manual*.

General Admin Permissions

Dashboards can be administered by any user with:

- permissions to configure dashboards
- permissions to administer a specific dashboard

  Note: Any individual who can configure an individual dashboards can delete it.
  - To configure dashboards, select **Configure** > **Dashboards** in the Portal menu.
  - If you cannot access this menu item, the menu item must be added to your menu profile. See “cxView User Administration” on page 184 in the “cxView User Administration” on page 184 page.

Access Permissions for Individual Dashboards

To view all contents of a dashboard, you must access to the Dashboard menu and permissions to view the specific dashboard and all components inside of it.

For example, if the dashboard contains scorecards, you must have permissions to view each scorecard, as well as access to the scorecard feature in the Portal (**Analyze** > **Scorecards**), which is governed by IBM Tealeaf cxView group permissions. The same applies to Report Builder and Top Mover components of dashboards.

The steps below define how to how to enable access permissions to an individual dashboard. The following permissions are also required:

- Access to the Dashboard menu, which is governed through IBM Tealeaf cxView groups. See “cxView User Administration” on page 184.
- Access to the individual dashboard. See steps below.
- Access to the menu item for viewing any dashboard component inside the report. For example, to see a scorecard component inside a dashboard, each user must have access to **Analyze Scorecards**, which is governed through IBM Tealeaf cxView groups. See “cxView User Administration” on page 184.
  - For more information on accessing Report Builder and Top Mover menu items, see “CX User Administration” on page 165.
- Access to the individual report components in the dashboard.
  - See “Scorecard Access Permissions” on page 429.
  - See “Report Access Permissions” on page 431.
  - See “Top Mover Access Permissions” on page 432.

Permissions to administer an individual dashboard are defined in the Configuring Dashboards page.

Note: The owner of a dashboard can always administer it.

1. To configure dashboards, select **Configure** > **Dashboards** in the Portal menu.
2. In the left panel, click **Dashboards**.
3. The list of available dashboards is displayed.
4. To administer a dashboard, select it and click the Pencil ( ) icon.
5. Click the Permissions tab.
   - In the left panel, you can configure the user groups that are permitted to administer the dashboard.
   - In the right panel, you can configure the user groups that are permitted to see the dashboard. Users in these groups can also save the dashboard as a different name. As the creator of the copied dashboard, the user is automatically an administrator for it.
6. To save changes, click **Save**.
   - See "Configuring Dashboards" in the *IBM Tealeaf cxView User Manual*.

Report Access Permissions

The Tealeaf Report Builder enables the configuration and display of event and ratio data with dimensional data optionally applied to segment the report.

- To access the Tealeaf Report Builder, select **Analyze** > **Report Builder** in the Portal menu.
- See "Tealeaf Report Builder" in the *IBM Tealeaf Reporting Guide*.

General Admin Permissions

To administer reports in general, a user account must include the menu item for the Report Builder in the menu profile.

See “CX User Administration” on page 165 in the “CX User Administration” on page 165 page.

Note: In the Tealeaf Report Builder, the user account `ADMIN` can administer any report.

Tealeaf provides a report management facility where you can set permissions, remove reports, and manage the schedules for generating reports.

- See "Report Manager" in the *IBM Tealeaf Reporting Guide*.

Access Permissions for Individual Reports

In the Report Builder, individual reports can be configured to permit administration privileges to selected user groups.

Note:

1. The owner of a report can always administer it.
2. Reports provided by Tealeaf cannot be edited. However, they can be saved as a separate reports and then modified.

1. From the Portal menu, select **Analyze** > **Report Builder**.
2. The Report Builder is displayed.

3. In the toolbar, click the Open ( ) icon.
4. Select the report to administer. Click **Open**.
5. In the toolbar, click the Permissions ( ) icon.
6. The Report Access Permissions dialog is displayed.
   - In the left panel, you can select the user groups that are permitted to administer the report.
   - In the left panel, you can select the user groups that are permitted to use the report.
   - See "Tealeaf Report Builder" in the *IBM Tealeaf Reporting Guide*.
7. To save changes, click **Save**.

Top Mover Access Permissions

Top Movers reports enable you to choose to record and report on deviations in values for event and dimension data. These reports allow tracking of the most active data pertaining to your web application.

The same access and permission methods for the Tealeaf Report Builder are also applicable to Top Mover reports.

- To access the Top Movers report, select **Analyze** > **Top Movers** in the Portal menu.
- See "Analyzing Top Movers" in the *IBM Tealeaf Reporting Guide*.

For more information on administration, see "Report Manager" in the *IBM Tealeaf Reporting Guide*.

Session Segment Access Permissions

Session segments can be created in either of the following ways:

1. automatically created based on the results of a Portal search
2. imported through the Segment Builder Service

Note: The Segment Builder Service is a component of IBM Tealeaf cxConnect for Data Analysis, a separately licensable component of the IBM Tealeaf CX platform. please contact your IBM Tealeaf representative.

- See "Segment Builder Service" in the *IBM Tealeaf cxConnect for Web Analytics Administration Manual*.

General Admin Permissions

To administer session segments, you must have access to the Session Segments page.

- To access the Session Segments page, select **Analyze** > **Segments** > **Manage Session Segments**.
- This item must be in your menu profile. See “CX User Administration” on page 165 in the “CX User Administration” on page 165 page.

Note: The user account `ADMIN` can administer any session segment.

Access Permissions for Individual Segments

Individual users can share session segments that they have created with other users that can access the Session Segments page. Segments that are shared can be used and deleted by all users.

1. In the Portal menu, select **Analyze** > **Segments** > **Manage Session Segments**.
2. Right-click the segment to share and select **Share**.
3. The segment is now available to all users.
   - To stop sharing, right-click the segment to share and select **Stop Sharing**.
   - See "Managing Session Segments" in the *IBM Tealeaf cxResults User Manual*.

Chapter 19. Managing User Agents

Through Tealeaf, you can monitor the different types of user agents that contact your web application. Through a provided set of data objects, Tealeaf can identify the type of traffic that is requesting resources from your web application and then monitor counts and other information that pertains to the type of user agent. These objects can be used as the source data for developing a useful set of user agent-related reports, enabling you to closely monitor the composition of traffic to your site.

Overview of User Agents in Tealeaf

The term **user agent** identifies a broad class of entities that can communicate with web servers across the Internet. This broad class includes desktop browsers, bots, and mobile devices, among others.

With each request to a web server, most devices include a string that uniquely identifies the device. This string can then be compared a public standard to ascertain other characteristics of the device that submitted the string.

This section provides more information on how user agents are tracked in Tealeaf.

How user agent information is tracked

The following steps outline the basic method that is used by Tealeaf to track user agent information.

1. **Client sends user agent string:** When a request is submitted to a web server, the client (browser or native application) submits a string that contains the publicly identifiable user agent string in the following variable, called a **header**, in the request. Below is an example user agent string:

   ```
   User-Agent: Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 5.1; Trident/4.0;
   .NET CLR 1.1.4322; .NET CLR 2.0.50727; .NET CLR 3.0.04506.30; .NET CLR
   3.0.04506.648; .NET CLR 3.0.4506.2152; .NET CLR 3.5.30729; .NET4.0C)
   ```

   Note: Depending on the type of client, the user agent string may be delivered in a different manner. See “User agent strings by client type” on page 436 below.
2. **Perform lookup of the user agent string against public standards:** When the user agent string is detected, Tealeaf performs a lookup against the applicable public standards:
   - Browscap for fixed desktop browsers
   - WURFL for mobile web browsers
   - Self-reporting from the Tealeaf client frameworks for mobile native applications
   - This step is performed by the Tealeaf Reference session agent.
     - This processing step in the Windows pipeline is included in the default pipeline. No additional configuration is required.
     - See "Tealeaf Reference Session Agent" in the *IBM Tealeaf CX Configuration Manual*.
3. **Populate request variables with data:** If a match between the submitted user agent string is found in one of the applicable standards, additional information

© Copyright IBM Corp. 1999, 2014

contained in the public standard entry for the user agent is populated in custom variables stored in the request by Tealeaf. See “How user agent information is stored” on page 437 below.

4. **Use request data:** Various Tealeaf components can reference user agent information for search, replay, and reporting purposes, which are explained at the end of this scenario, after data objects have been configured.

User agent strings by client type

Depending on the type of client device used by the visitor to experience your web application, its user agent information may be forwarded in a different manner for verification by the listed public standard.

Table 41. User agent strings by client type

Client Type	Source of user agent string	Standard used
fixed desktop browser	User agent string is submitted by default by the browser	browscap
Android mobile native application	Tealeaf client framework submits user agent string **Note:** Capture of Android-based mobile native applications requires the IBM Tealeaf CX license, the IBM Tealeaf CX Mobile license, and deployment of the Android Logging Framework. See "Tealeaf Android Logging Framework Reference Guide" in the *IBM Tealeaf Android Logging Framework Reference Guide*. **Note:** Since mobile native applications submit a user agent string only if specified by the developer, the Android Logging Framework submits a string as well as a set of properties in the `HTTP_X_TEALEAF_PROPERTY` header. • For more information on the submitted properties, see "Tealeaf Reference Session Agent" in the *IBM Tealeaf CX Configuration Manual*.	WURFL

Table 41. User agent strings by client type (continued)

Client Type	Source of user agent string	Standard used
iOS mobile native application	Tealeaf client framework submits user agent string. **Note:** Capture of IOS-based mobile native applications requires the IBM Tealeaf CX license, the IBM Tealeaf CX Mobile license, and deployment of the IOS Logging Framework. See "Tealeaf iOS Logging Framework Reference Guide" in the *IBM Tealeaf iOS Logging Framework Reference Guide*. **Note:** Since mobile native applications submit a user agent string only if specified by the developer, the iOS Logging Framework submits a string as well as a set of properties in the `HTTP_X_TEALEAF_PROPERTY` header. • For more information on the submitted properties, see "Tealeaf Reference Session Agent" in the *IBM Tealeaf CX Configuration Manual*.	WURFL
IBM Tealeaf CX UI Capture for AJAX fixed desktop browser	User agent string is submitted by default by browser **Note:** Capture of AJAX-based web applications requires the IBM Tealeaf CX license and deployment of the IBM Tealeaf CX UI Capture for AJAX client framework. See "UI Capture for AJAX Guide" in the *IBM Tealeaf UI Capture for AJAX Guide*.	browscap
IBM Tealeaf CX UI Capture for AJAX mobile web browser	User agent string is submitted by default by the mobile web browser. **Note:** Capture of AJAX-based mobile web applications requires the IBM Tealeaf CX license, the IBM Tealeaf CX Mobile license, and deployment of the IBM Tealeaf CX UI Capture for AJAX client framework. See "UI Capture for AJAX Guide" in the *IBM Tealeaf UI Capture for AJAX Guide*. **Note:** Since the user agent information is submitted by the browser, mobile web applications captured by UI Capture do not include the additional properties that are submitted in the `HTTP_X_TEALEAF_PROPERTY` header for mobile native applications.	WURFL

How user agent information is stored

When a match is found for the user agent string submitted by the client, the following information is inserted into the `[ExtendedUserAgent]` section of the request.

Common values

The following values are populated for a matched user agent:

```
[ExtendedUserAgent]
TLT_BROWSER=Firefox
TLT_BROWSER_VERSION=Firefox3.6
TLT_BROWSER_PLATFORM=Win XP
TLT_TRAFFIC_TYPE=BROWSER
TLT_BROWSER_JAVASCRIPT=true
TLT_BROWSER_COOKIES=true
```

TLT_TRAFFIC_TYPE request variable

As part of its processing, the Tealeaf Reference session agent inserts the `TLT_TRAFFIC_TYPE` request variable, as in the following example:

```
TLT_TRAFFIC_TYPE=BOT
```

These values populate the `Traffic Type` dimension provided by Tealeaf.

- See "TEM Dimensions Tab" in the *IBM Tealeaf Event Manager Manual*.

Special Values

For the following values for `TLT_TRAFFIC_TYPE`, additional request variables may be inserted.

Value	Description
MOBILE	`TLT_TRAFFIC_TYPE=MOBILE` indicates a session gathered from a mobile web browser through IBM Tealeaf CX UI Capture for AJAX. • See “Mobile Web sessions” below. • For more information on UI Capture, see "UI Capture FAQ" in the *IBM Tealeaf UI Capture for AJAX FAQ*.
MOBILE_APP	`TLT_TRAFFIC_TYPE=MOBILE_APP` indicates a session gathered from a mobile native application monitored by one of the Tealeaf client frameworks. • See “Mobile native app sessions” below. • See "Tealeaf Android Logging Framework Reference Guide" in the *IBM Tealeaf Android Logging Framework Reference Guide*. • See "Tealeaf iOS Logging Framework Reference Guide" in the *IBM Tealeaf iOS Logging Framework Reference Guide*.

Mobile Web sessions

If `TLT_TRAFFIC_TYPE=MOBILE`, then the following fields may be inserted and populated for the Mobile Web user agent:

```
TLT_SCREEN_HEIGHT
TLT_SCREEN_WIDTH
TLT_COLOR_DEPTH
TLT_PICTURE_SUPPORT
TLT_VIDEO_SUPPORT
TLT_STREAMING_SUPPORT
```

Mobile native app sessions

User agent information for a mobile native application is submitted at the discretion of the application developer. To enable normalized user agent information, each Tealeaf client framework submits a specific header, which enables detection of user agent information. When this header and the associated

properties are detected in a hit, the following values are populated in the [ExtendedUserAgent] based on gathered and submitted data from the client framework.

- For more information about user agent detection for mobile native applications, see "Tealeaf Reference Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

```
[ExtendedUserAgent]
TLT_BROWSER=StraussAndPlesser Native
TLT_BROWSER_VERSION=8.5
TLT_BROWSER_PLATFORM=Android
TLT_BRAND=Asus
TLT_MODEL=Asus Eee Pad Transformer TF101
TLT_SCREEN_HEIGHT=800
TLT_SCREEN_WIDTH=1280
TLT_COLOR_DEPTH=65536
```

Enabling extended user agent detection

Before you begin, you must configure the Tealeaf Reference session agent to enable extended user agent parsing. When enabled, captured user agent information is posted into the request in the [ExtendedUserAgent] section by the Tealeaf Reference session agent.

- See "Tealeaf Reference Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

Mobile User Agents and bots

In particular, Tealeaf provides mechanisms for monitoring traffic from mobile devices and from bots.

The number of mobile web visitors continues to expand at a rapid rate. Depending on your web application, these visitors may have special requirements and customer issues, so the ability to track their activities independent of desktop and other types of visitors is essential in providing good service to this market.

- The IBM Tealeaf CX Mobile module enables insightful tracking and useful monitoring of mobile visitors and their devices. See "Tealeaf CX Mobile Administration Manual" in the *IBM Tealeaf CX Mobile Administration Manual*.

Short for "robot," a **bot** is a program designed to crawl pages of a website. A bot can be used for indexing a website for search engines, accumulating price information, and other organizational purposes. Tealeaf enables the monitoring of bot activities on your website. This type of activity can be tracked easily through RTV and Portal searches.

- See "bot Detection" in the *IBM Tealeaf cxImpact Administration Manual*.

Public Standards

To monitor user agents, Tealeaf utilizes two publicly maintained standards. These standards are individual files containing all known user agents and mobile devices.

- BrowsCap.csv is the standard for fixed user agents.
- WURFL.csv is the standard for mobile devices.

You can also customize your own set of user agents and apply it to your Tealeaf installation. Through UserCap.csv, you can provide configuration information for fixed and mobile user agents. These settings can be used to add new user agents to

the standards listed above. For more information on how Tealeaf uses these configuration files, see "Configuring User Agent Detection" in the *IBM Tealeaf cxImpact Administration Manual*.

Note: These files should be updated on a regular basis. See "Maintaining the CX System" in the *IBM Tealeaf CX Installation Manual*.

User Agent Tools

To facilitate easy management of these configuration files, Tealeaf provides a set of external tools for converting the files to a format that Tealeaf can use and for reviewing the uploaded files. See "User Agent Tools" in the *IBM Tealeaf cxImpact Administration Manual*.

User Agent Events

You can monitor and report on the user agents visiting your site. Tealeaf provides a set of default data objects that you can enable and use to develop the events required to meet your reporting needs.

- See "Configuring User Agent Events" in the *IBM Tealeaf cxImpact Administration Manual*.

Tutorial

For more information on how user agent information is used within Tealeaf, including steps to configure event objects to track user agent information, see "E2E Scenario - Tracking User Agent Information in Tealeaf" in the *IBM Tealeaf cxImpact User Manual*.

Chapter 20. Configuring User Agent Detection

This section describes how to acquire and maintain the configuration files necessary for parsing of fixed and mobile user agents and how to prepare those files for use in the IBM Tealeaf CX platform.

Note: User agent detection must be enabled through the Tealeaf Reference session agent. Steps are provided at the conclusion of this section.

Types of Files

There are several files needed by Tealeaf to collect and display various user agent information for reporting.

1. `WURFL.csv` - This file contains the library of user agent information for mobile devices. This file is maintained by a third-party and should be updated in your Tealeaf solution on a regular basis. See "`WURFL.csv`."
2. `browscap.csv` - A publicly available standard for user agents on the web. This file is maintained by a third-party and should be updated in your Tealeaf solution on a regular basis. See "`browscap.csv`" on page 442.
3. `UserCap.csv` - This file is created and maintained by Tealeaf in a format similar to the above `browscap.csv`. You can use this file to overwrite the values for specific user agents in the `browscap.csv` standard.
   - This file may be used to insert user agent values that are not present in the public standard.
   - See "`UserCap.csv`" on page 442.
4. `UserSupplement.csv` - This file is created and maintained by Tealeaf in a format similar to the above `browscap.csv`. You can use this file to add user agent values that do not appear in either public standard. See "`UserSupplement.csv`" on page 443.

WURFL.csv

The `WURFL.xml` public standard contains all the user agent information for known mobile devices.

Note: Management of mobile user agents requires the IBM Tealeaf CX Mobile module, a separately licensable component of the IBM Tealeaf CX platform. please contact your IBM Tealeaf representative.

- For more information on enabling, see "Overview of CX Mobile" in the *IBM Tealeaf CX Mobile User Manual*.
- For more information on data acquired from mobile devices, see "Overview of CX Mobile" in the *IBM Tealeaf CX Mobile User Manual*.

Periodically, you should update your local version of WURFL from the public standard.

Note: Do not make changes to this standard-based file, which is overwritten whenever an update is made from the public standard. Apply new entries or update entries in the `UserSupplement.csv` and `UserCap.csv` files, respectively.

© Copyright IBM Corp. 1999, 2014

Note: WURFL.XML should be downloaded and updated in your Tealeaf system on a regular basis. See "Maintaining the CX System" in the *IBM Tealeaf CX Installation Manual*.

When it is downloaded, WURFL.xml must be converted from its source .xml file to the .csv format, which is required for Tealeaf use, using the WURFL2CSV.exe utility.

- Before it is used, the Tealeaf software verifies that this file is in the proper format.
- For more information on converting WURFL.xml, see "User Agent Tools" in the *IBM Tealeaf cxImpact Administration Manual*.

browscap.csv

This file contains a list of all known user agents and system information that is related to them. This file is maintained by Gary Keith and can be collected from his website.

This file is available in various formats but only the .csv file format is supported by Tealeaf. This file is verified by Tealeaf to be in the original format.

Note: Do not change this standard-based file, which is overwritten whenever an update is made from the public standard. Apply new entries or update entries in the UserSupplement.csv and UserCap.csv files, respectively.

Note: browscap.csv should be downloaded and updated in your Tealeaf system regularly. See "Maintaining the CX System" in the *IBM Tealeaf CX Installation Manual*.

UserCap.csv

In some cases, you may need to overwrite the user agent strings and their corresponding data that are available in the public standards. For example, you may need to correct errors that your enterprise has detected in the public standards with more accurate information.

Tealeaf provides the UserCap.csv file, which can be used to override values in the public standards. By gathering the user agent string value from the source file, you can then create a new entry in the UserCap.csv file through the UserAgentRevealer.exe utility. During evaluation of user agent strings, the Tealeaf Reference session agent checks UserCap.csv to see if it contains replacement values for any detected user agents. If so, the replacement values are used instead.

Note: Configurations in UserCap.csv replace corresponding entries in browscap.csv or WURFL.csv. If you wish to add values that do not appear in these files, use UserSupplement.csv. For more information on replacing values by configuration file, see "UserSupplement.csv" on page 443.

Note: Be careful in creating replacement user agent strings. Since UserCap.csv takes precedence over the public standards at runtime, it is possible for the public standards to become updated without having those updates reflected in the user agent data consumed by Tealeaf, which is using override values. The contents of UserCap.csv should be reviewed whenever updates are made to the public standards in use by your Tealeaf implementation.

An example version of UserCap.csv is included during installation. You may configure entries in this file using the UserAgentRevealer.exe utility. See "Reviewing Configuration Files."

UserSupplement.csv

If you need to add custom user agent strings for fixed or mobile user agents, you can add the configurations to this file, which is used for adding new definitions not covered in browscap.csv and WURFL.csv definitions.

This mechanism allows you to include anomalous or unreported data for purposes of reporting. For example, you can use this file to include a single value for reporting purposes without updating all configuration files.

Note: Configurations in UserSupplement.csv do not override any corresponding entries in browscap.csv or WURFL.csv, which are public standards. For more information on replacing those values, see "UserCap.csv" on page 442.

An example version of UserSupplement.csv is included during installation. You may configure entries in this file using the UserAgentRevealer.exe utility. See "Reviewing Configuration Files."

Creating UserSupplement.csv

By default, a version of UserSupplement.csv is not provided with your installation. If you need to create a version of UserSupplement.csv, please complete the following steps.

1. Login to one of the Procesing Servers (Canisters) as an administrator.
2. Navigate to the following directory:
   <Tealeaf_install_directory>\System
3. Create a copy of usercap.csv.
4. Rename the copied file: UserSupplement.csv.
5. Edit the file in Microsoft Excel.
6. Remove all data lines in the file so that only the header remains.
7. Add your entries to the file using UserAgentRevealer.exe. See "Reviewing configuration information" on page 445.
8. Save the file and store it in the same directory.
9. This file must be copied to all Canisters that process hits.
10. Restart the Transport Service on each Canister.

Reviewing Configuration Files

Tealeaf includes the User Agent Revealer utility (UserAgentRevealer.exe) for reviewing the contents of the user agent configuration files, which can grow quite large. You can enter a user agent string into the utility, which then displays how the Tealeaf installation reports the string, after comparing the provided values against all available configuration files.

- This tool is used primarily for debugging purposes. You must run the tool from the machine where it is stored.

  Note: This tool requires access to browscap.csv, which must be installed in the system directory inside the Tealeaf install directory.

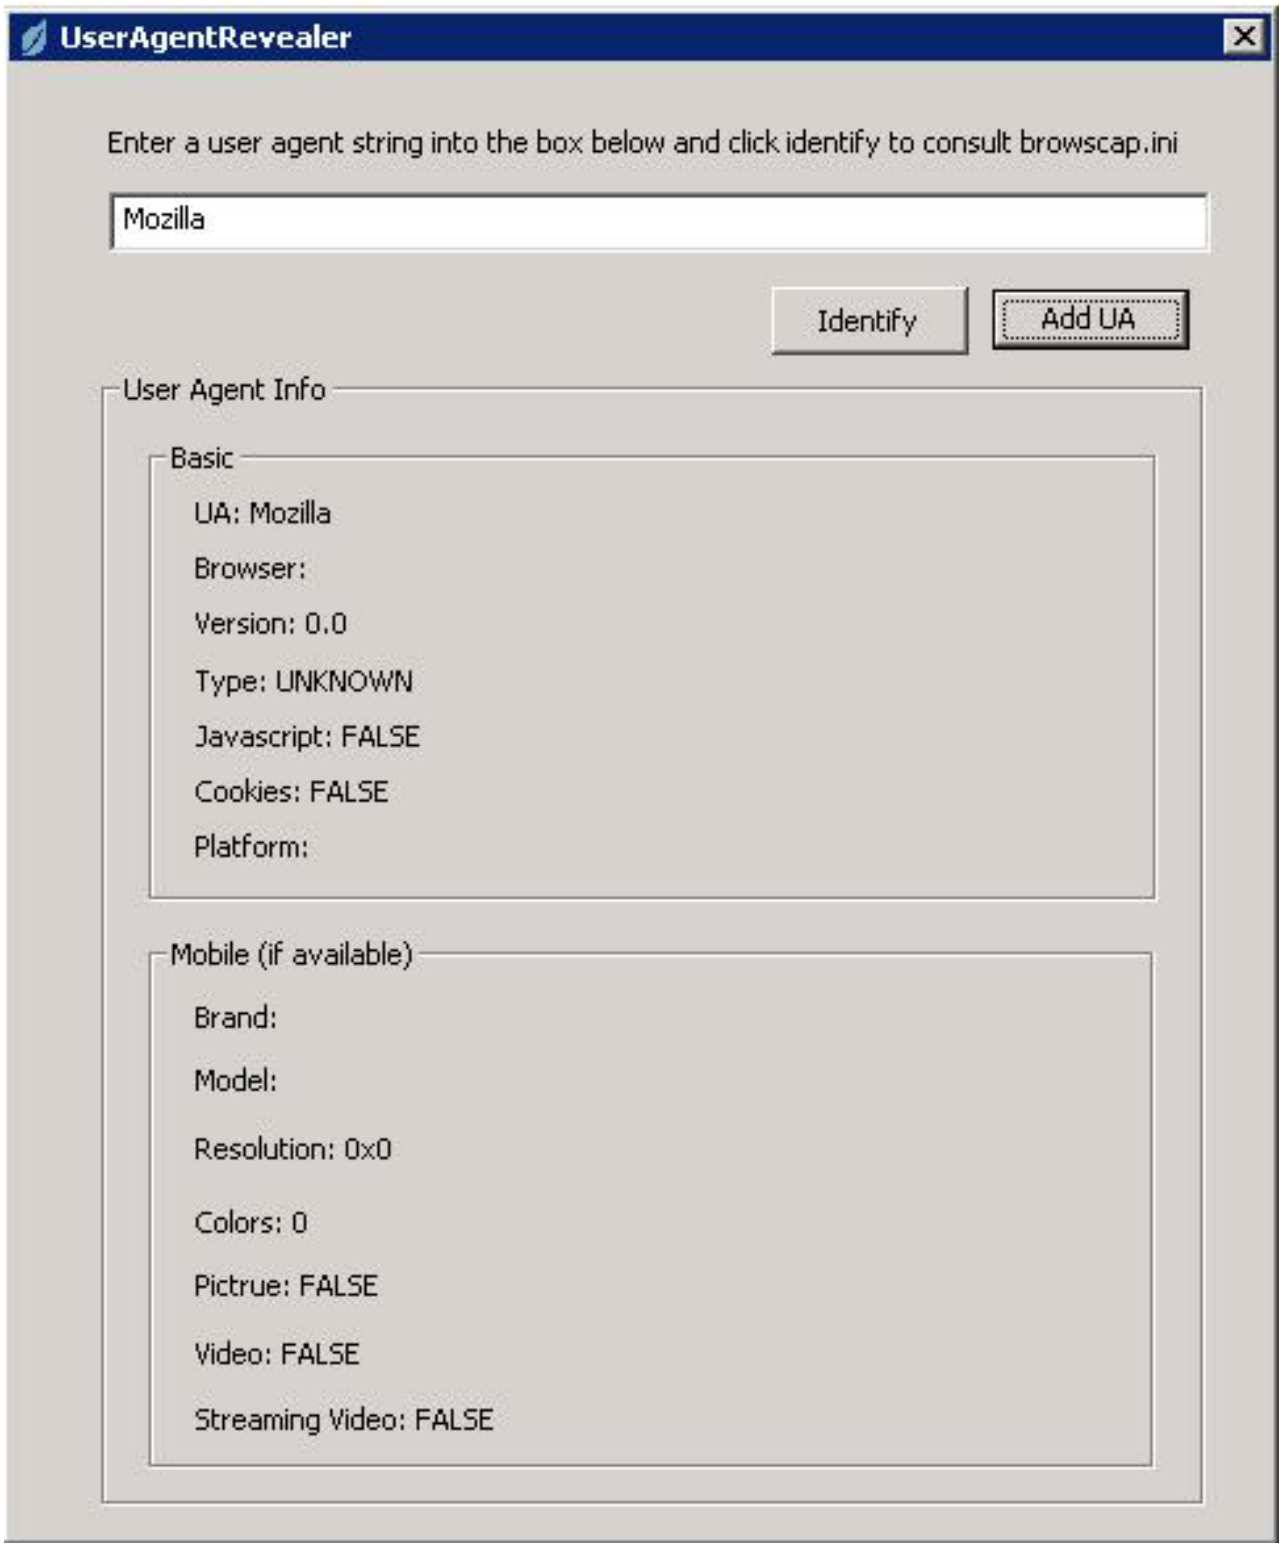

Figure 103. User Agent Revealer

Reviewing configuration information

1. In the Tealeaf install directory, double-click `UserAgentRevealer.exe`.
2. Paste a user agent string into the textbox. Click **Identify**.
3. Any available information for the user agent is displayed.
   - If no match is found, you may create a new entry in `UserSupplement.csv`.
   - If a match is found with incorrect data in it, you may create an entry in `UserCap.csv` to replace it.
   - See "Editing or creating user agent entries."

Editing or creating user agent entries

You can edit found entries or add new ones through `UserAgentRevealer.exe`.

1. To edit an identified user agent entry or to create a new one, click **Add UA**.

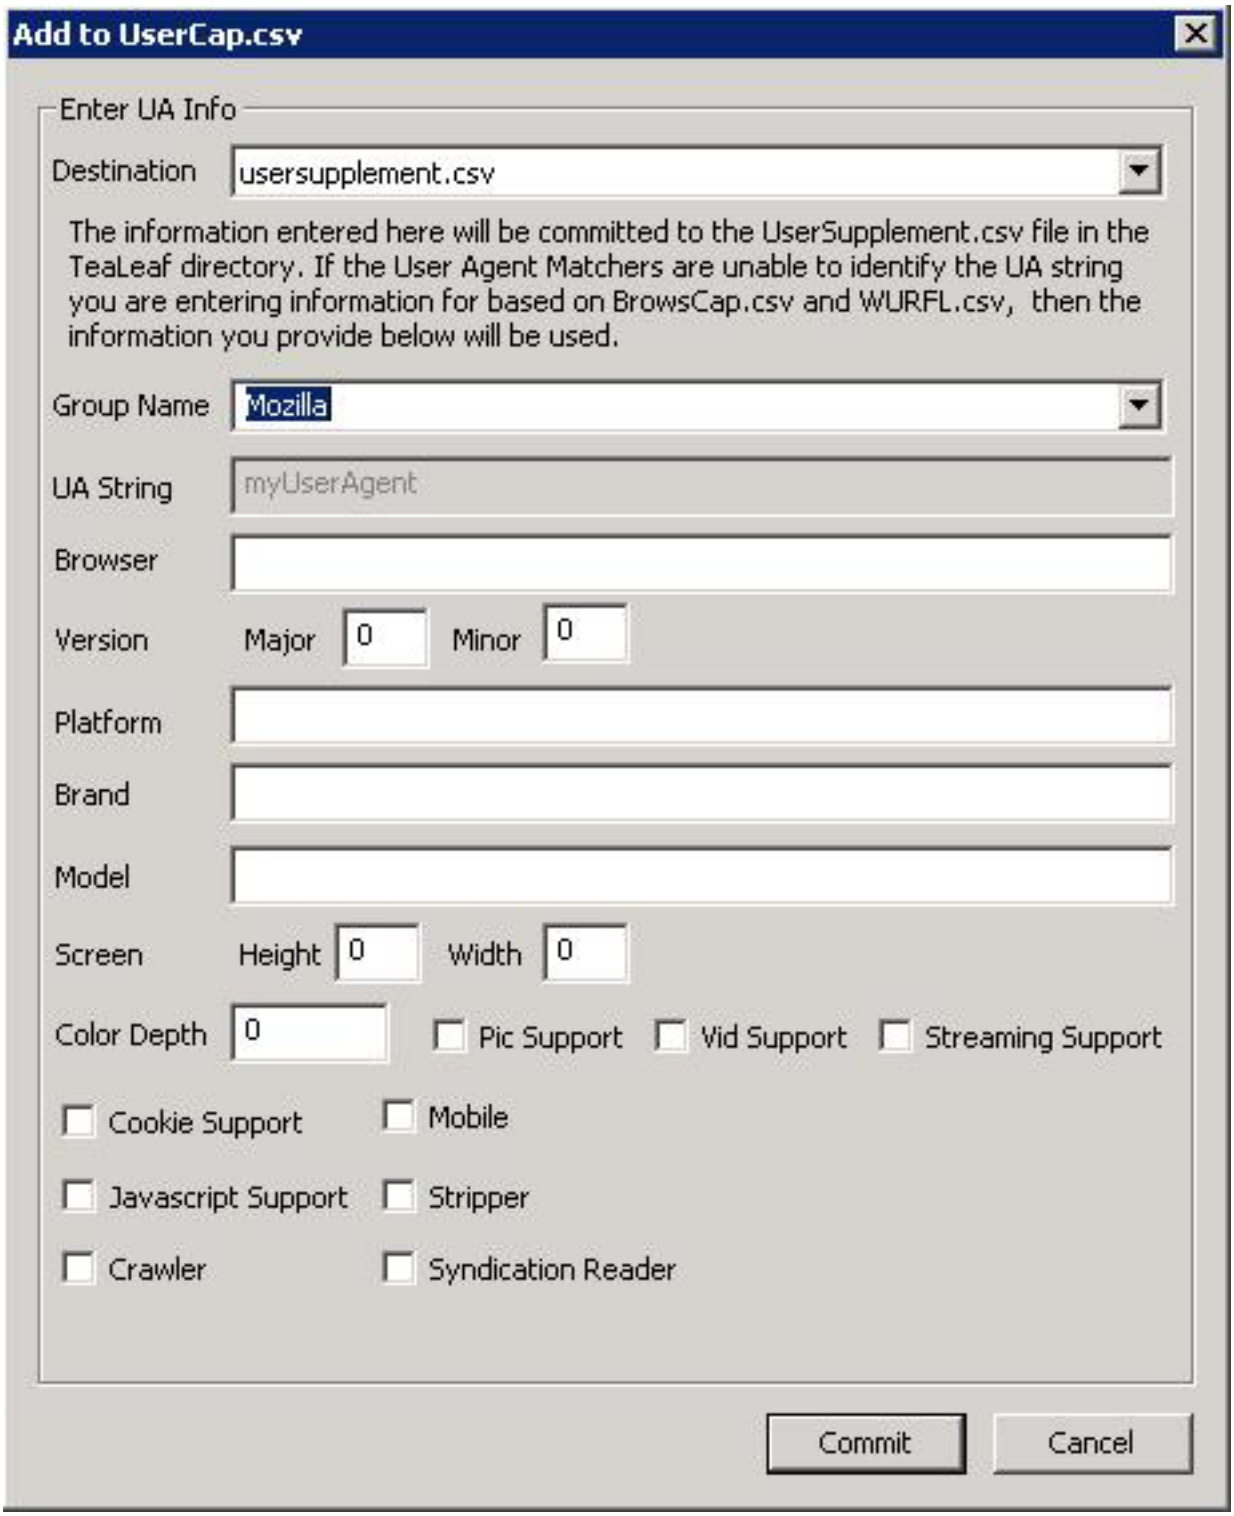

Figure 104. Adding User Agent Entry to usersupplement.csv

2. Select the appropriate Destination:
   - `usersupplement.csv` - If the user agent entry cannot be identified, the Type value is set to `UNKNOWN`. In these instances, you can add the user agent string to the `usersupplement.csv` file with values that you define.

- `usercap.csv` - If the user agent entry contains errors or values that you would like to change, you can create an entry with the overriding data in `usercap.csv`. These values are used instead of the values listed in the applicable public standard (`browscap.csv` or `WURFL.csv`).

3. Fill out all relevant fields.

   Note: All information entered into the fields is case-sensitive. Double quotes cannot be used and are stripped from the input.

   - You may enter text into the Group field to create a new Group name, if needed.
   - See "Default User Agent Fields."
4. To commit the changes, click **Commit**.
5. The selected Destination file is updated and saved.
6. When Tealeaf next refreshes its cache from the Destination file, the user agent information is applied to all subsequent hits.

Default User Agent Fields

The following fields are available by default for user agent tracking:

Field Description

`Destination`
User-defined `.csv` where the new or edited entry is saved.
- `usersupplement.csv` - for new user agent entries
- `usercap.csv` - for edited user agent entries

`Group Name`
User agent group name to which the new agent belongs

`UA String`
String you entered in the Identify textbox

`Browser`
Browser type

`Version`
Version number, including major and minor version indicators. For example, the entries for `IE 6.0` in these fields would be `6` and `0`

`Platform`
Operating system

`Brand` Brand name for the user agent

`Model` Model name or number

`Screen` Screen size in `Height` and `Width` in pixels

`Color Depth`
The number of bits of color depth. For example, if the user agent has 32-bit color depth, enter `32`.

`Pic Support`
Select if the user agent supports display of picture formats

`Vid Support`
Select if the user agent supports display of video formats

`Streaming Support`
Select if the user agent supports display of streaming formats

Cookie Support
Select if the user agent supports cookies

Javascript Support
Select if the user agent supports Javascript

Crawler
Select if the user agent is a bot

Mobile Select if the user agent is a mobile device

Stripper
Select if the user agent is a stripper (a bot that scrapes websites for content)

Syndication Reader
Select if the user agent is a reader, such as an RSS feed

Enable extended user agent detection

After the configuration files were downloaded and prepared for use, you can configure the Tealeaf Reference session agent to use them for user agent detection, if you did not do so already.

Note: By default, extended user agent parsing is enabled.

See "Tealeaf Reference Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

Chapter 21. Configuring User Agent Events

To use the data that is injected into the [ExtendedUserAgent] section by the Tealeaf SessionAgentTLTRef, a Tealeaf administrator must create events that create charts for this data.

User Agent Evaluation Logic

Self-reporting user agents can be identified through the user-agent string value in the request. Tealeaf matches the values reported in this field to the values in the UserAgent field in the public standard Browscap.csv.

If matches are found, then the following fields in the agent's Browscap.csv record are evaluated:

- Crawler
- SyndicationReader
- MobileDevice

Possible values for the above fields are the following:

- TRUE
- FALSE

The evaluation logic is as follows:

- If MobileDevice=true, then the user agent is identified as a mobile device.
- If Crawler=true AND MobileDevice=true, then the user agent is identified as a mobile_bot by Tealeaf.
- If Crawler=true OR SyndicationReader=true, then the user agent is identified as a bot by Tealeaf.
- Otherwise, the user agent is identified as browser.
- If the field is set to UNKNOWN, no data could be found for the UserAgent in any of the files. Session data from these user agents should be examined.

Note: The above information also applies to entries in the UserCap.csv file.

User Agent Event Objects

Tealeaf provides some objects to monitor user agents detected in the capture stream. You can create additional data objects to track all user agent information.

- See "E2E Scenario - Tracking User Agent Information in Tealeaf" in the *IBM Tealeaf cxImpact User Manual*.

Locating a Missing User Agent

Some User Agents may not be included in either the BrowsCap or WURFL user agent file.

To check if the identifying string for a user agent that you are trying to track is in the files, you can use the UserAgentRevealer utility. See "User Agent Tools" in the *IBM Tealeaf cxImpact Administration Manual*.

© Copyright IBM Corp. 1999, 2014

If your tested string is not found in either file, custom user agent strings can be created and integrated with Tealeaf through the `UserCap.csv` file. See "Configuring User Agent Detection" in the *IBM Tealeaf cxImpact Administration Manual*.

Adding Values for Reporting

For reporting purposes, Tealeaf provides the `Traffic Type` dimension, which can be used to segment your reports based on the type of user agent. It pulls values from the `Traffic Type` system hit attribute.

See "Tealeaf Report Builder" in the *IBM Tealeaf Reporting Guide*.

To create reports that operate on the data presented by BrowsCap and WURFL, you must acquire the captured values and store them in dimensions you create.

The best method for capturing this data is to create dimensions to store it, if necessary. When you enable logging on those dimensions, you can later download the dimension values, normalize them as needed, in a text file. This file can then be imported as a whitelist set of values back to the dimension through the Tealeaf Event Manager.

- See "User Agent Tools" in the *IBM Tealeaf cxImpact Administration Manual*.
- These values can then be added to dimensions you create in the Event Manager. See "TEM Dimensions Tab" in the *IBM Tealeaf Event Manager Manual*.

Chapter 22. User Agent Tools

The following external utilities can be used to monitor and update user agent configuration data required by Tealeaf.

Converting WURFL

You can use the `WURFL2CSV.exe` utility to convert the downloaded source `.xml` for the WURFL file to the `.csv` format, which is required by Tealeaf.

- For more information on where to acquire WURFL, see "Maintaining the CX System" in the *IBM Tealeaf CX Installation Manual*.

  Note: Conversion using `WURFL2CSV.exe` requires the IBM Tealeaf CX Mobile module, a separately licensable component of the IBM Tealeaf CX platform. please contact your IBM Tealeaf representative.

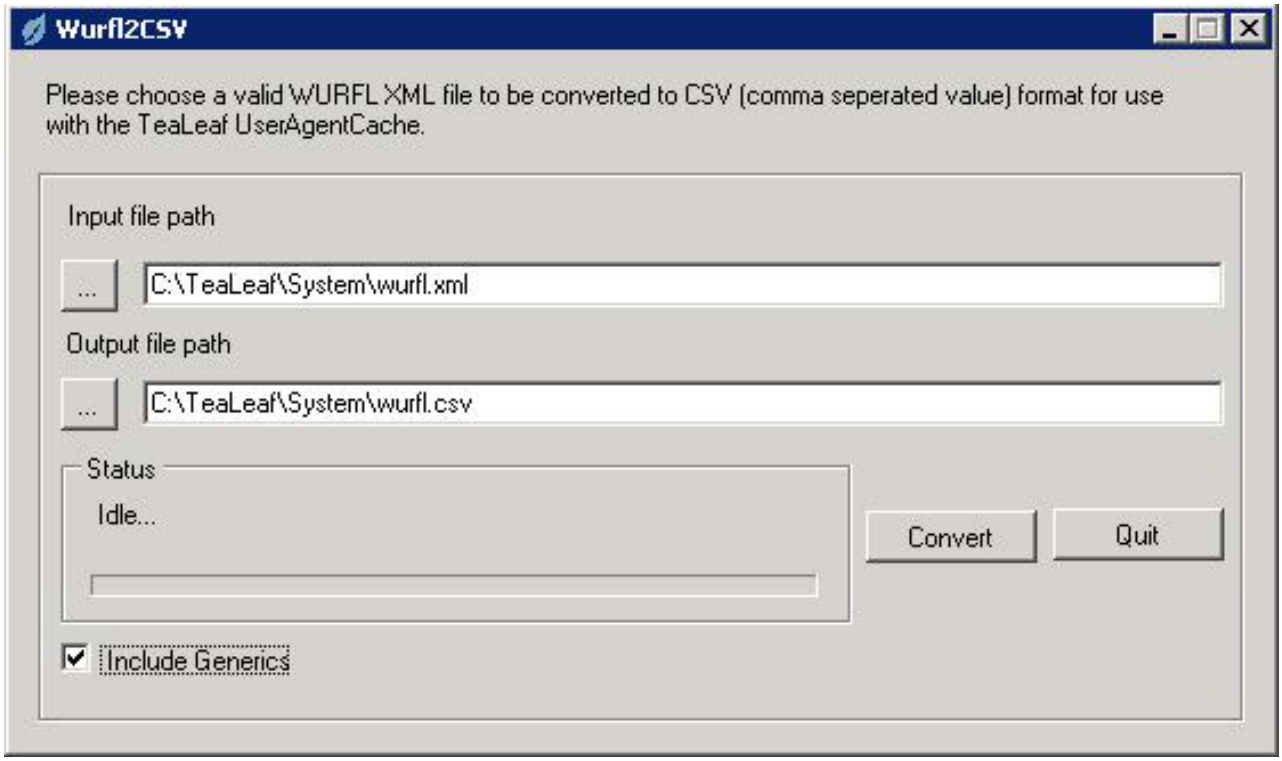

Figure 105. WURFL2CSV.exe

1. Double-click `WURFL2CSV.exe`.
2. Select the downloaded `WURFL.xml` as the input file.
3. Provide an output filename or use the default `WURFL.csv` filename. Store the file in the location suggested by the application.
4. If you wish to add generic versions of mobile user agents to the `.csv` version.
   - Generics are used as a fallback entry for partial matches. For example, if `iPhone 10` is detected in the capture stream and no exact match exists in `WURFL.csv`, a generic user agent string (`iPhone`) can be used instead.

© Copyright IBM Corp. 1999, 2014

- By default, use of generics is enabled. See "Tealeaf Reference Session Agent" in the *IBM Tealeaf CX Configuration Manual*.

5. To generate the `.csv` file, click **Convert**.
6. The file is generated and stored in the specified location.

For more information on WURFL, see "Configuring User Agent Detection" in the *IBM Tealeaf cxImpact Administration Manual*. For more information on where to acquire WURFL, see "Maintaining the CX System" in the *IBM Tealeaf CX Installation Manual*.

Creating User Agent Event Value Lists

Note: In Release 8.5 and earlier, Tealeaf provided the `UAValueListCreator.exe` utility for the creation of value lists. This utility has been superseded by gathering detected values from the logs for the dimensions where they are stored This section describes the new process.

In order to begin gathering user agent information, you must add or configure the corresponding dimensions to gather and store values. These functions are performed through the Event Manager.

- You can use this method for capturing user agent information from the `Browscap.csv` standard for fixed user agents and from the `WURFL` standard for mobile user agents.
- For more information on the configuration files, see "Configuring User Agent Detection" in the *IBM Tealeaf cxImpact Administration Manual*.

For step-by-step instructions on how to create data objects to populate dimension value lists, see "E2E Scenario - Tracking User Agent Information in Tealeaf" in the *IBM Tealeaf cxImpact User Manual*.

Reviewing configuration files

`UserAgentRevealer.exe` can be used to review the `.csv` configuration files for user agents. See "Configuring User Agent Detection" in the *IBM Tealeaf cxImpact Administration Manual*.

Chapter 23. bot Detection

Most commercial enterprise web applications are subjected to significant traffic from web crawling bots. Short for "robot," a bot is a software agent that navigates websites to extract information about them. Some bots are used by search engines for indexing pages.

Through its user agent detection capabilities, Tealeaf enables you to closely monitor activities on your web application that are identified as bot traffic.

Enabling bot Detection in the CX Platform

Before you begin, you must configure Tealeaf to detect user agents.

Please complete the following steps.

1. You must verify installation of the appropriate user agent detection files. For best results, you should update your local copies of these files from the available public standards. See "Configuring User Agent Detection" in the *IBM Tealeaf cxImpact Administration Manual*.
2. After you have acquired and installed the user agent files, you need to enable user agent detection through the Tealeaf Reference session agent, which must be included in your Windows processing pipeline.
   - For more information on enabling extended user agent detection, see "Tealeaf Reference Session Agent" in the *IBM Tealeaf CX Configuration Manual*.
   - For more information on configuring your pipeline, see Chapter 8, “TMS Pipeline Editor,” on page 307.

When the above steps are completed, Tealeaf can perform the basic detection of user agents, including bots, utilizing the public standards.

Detecting bots

Valid bots identify themselves as user agents of `bot` type and can be tracked through Tealeaf. For example, when bots are detected, you can identify the pages on your site that have been crawled by specific search engine bots such as googlebot, msnbot, and more.

Self-reporting bots can be identified through the `user-agent` string value in the request. Tealeaf can match the values reported in this field to the values in the `UserAgent` field in the public standard `Browscap.csv`. This user agent information is posted to the `[ExtendedUserAgent]` area of the request.

- See Chapter 21, “Configuring User Agent Events,” on page 449.

Optionally, you can configure the Tealeaf Reference Session Agent to overload the `TLT_APPLICATION_NAME` with the browser type extracted from the user agent cache.

- See "Tealeaf Reference Session AgentOverloadingApplicationReferenceField" in the *IBM Tealeaf CX Configuration Manual*.

© Copyright IBM Corp. 1999, 2014

Masked bots

Some bots do not report themselves as legitimate bots. These bots may attempt to crawl your application for illegitimate purposes or competitive advantages. Tealeaf does not provide direct mechanisms for identifying masked bot activity.

All known and self-reporting bots are maintained in `Browscap.csv`, which Tealeaf uses to identify user agents.

If you detect a user agent that is not properly captured by Tealeaf, you can add it to `UserCap.csv` for use by Tealeaf. See "User Agent Tools" in the *IBM Tealeaf cxImpact Administration Manual*.

Note: Identifying masked bots may require additional programming and customization. For more information, please contact Tealeaf Professional Services.

bot-Related Data Objects

Tealeaf provides the following bot-related data objects that can be used for monitoring and reporting on self-reporting bots.

Bot (T/F) System Session Attribute

Tealeaf provides a session attribute (`Bot T/F`), which is set to `true` if the session applies to a self-reporting bot.

This session attribute is used for the following purposes:

- In the session, this item is written into the `[canisterSummary]` section, which can be viewed through RTV. See “bot Detection in RTV.”
- For reporting purposes, you can use the Traffic Type dimension, one of the values of which is `BOT`. See “Traffic Type Dimension.”

bot Detection in RTV

Tealeaf's provided session attributes are written to the `STS` file associated with each session. Among the session attributes reported in the `[canisterSummary]` section, the following session attribute indicates that the user agent is a self-reporting bot.

```
TltStsIsBot=1
```

- See Chapter 21, “Configuring User Agent Events,” on page 449.
- These session attributes are available for review through RTV. See "RealiTea Viewer - Session Attributes" in the *IBM Tealeaf RealiTea Viewer User Manual*.
- Some of these fields are automatically indexed for search. See "Configuring CX Indexing" in the *IBM Tealeaf CX Configuration Manual*.

Traffic Type Dimension

Tealeaf provides a dimension to identify the type of user agent. For user agents that have been identified by Tealeaf as a bot, the dimension `Traffic Type` is recorded with a value of `BOT`. In the Tealeaf Report Builder, this dimension can be added to your reports to filter event and ratio metrics to display only bot-based values.

- For more information on adding dimensions to your report, see "Tealeaf Report Builder" in the *IBM Tealeaf Reporting Guide*.
- For more information on filtering the values of added dimensions, see "Tealeaf Report Builder" in the *IBM Tealeaf Reporting Guide*.

- For more information on using this dimension, see “Reporting on bots.”
- For more information on creating or editing dimensions, see "TEM Dimensions Tab" in the *IBM Tealeaf Event Manager Manual*.

Searching for bot sessions

This section describes how to search for bot sessions.

Searching in RTV

In the IBM Tealeaf CX RealiTea Viewer, you can search for bot traffic.

To specify the search:

1. In RTV, click the Search button in the RTV toolbar.
2. At the bottom of the Search Builder, events available for search are listed as folders, with report groups and dimensions listed beneath them. Find an event that occurs in every session and contains the Traffic Type dimension.

   Note: You may have to create this event with `Traffic Type` as an included dimension. See "TEM Events Tab" in the *IBM Tealeaf Event Manager Manual*.
   - See "RealiTea Viewer - Session Search and Subsearch" in the *IBM Tealeaf RealiTea Viewer User Manual*.
3. Select the `Traffic Type` dimension.
4. In the right side of the panel, for the dimensional value, enter `BOT`.
5. Specify the other parameters of your search.
6. To begin the search, click **Search**.
   - See "RealiTea Viewer - Session Search and Subsearch" in the *IBM Tealeaf RealiTea Viewer User Manual*.

Searching through the Portal

Through the Tealeaf Portal, you can search for and replay sessions from bots.

1. In the Portal, select **Search** > **Completed Sessions**.
2. In the Search Fields panel, remove all default options.
3. In the Basic Search Fields panel, click **Event**.
4. In the Event Values area, click **<Select an event**. Select an event that contains the Traffic Type dimension.
5. In the Event Values area, click **<Any Dimension**. Select `Traffic Type`.
6. In the Search Scope drop-down, select `OR - Any Session`.
7. Specify the other parameters of your search.
8. To begin the search, click **Search**.

To replay a found search, click the results. Then, click the Replay icon for a specific session.

- To replay through Browser-Based Replay, click **Browser**. See "CX Browser Based Replay" in the *IBM Tealeaf cxImpact User Manual*.

Reporting on bots

For most reporting situations, you can use the Traffic Type dimension to segment your report to display only bot-related activities.

- You may add this dimension to your X-axis to show bot traffic side-by-side with other types of traffic
- You may add this dimension to the Y-axis to create stacked charts showing how bot traffic compares to other types of traffic in metric calculation.
- You may add this dimension as a Segment to show bot traffic only.
- See “Traffic Type Dimension” on page 454.

Chapter 24. Tealeaf Reference Information

Use the links below to access reference information on file formats, common standards, and more in use by the Tealeaf system.

Sections

- "Tealeaf Encryption Standards" on page 458
- "Tealeaf File Extensions"

Tealeaf File Extensions

Tealeaf solutions use a variety of common file types, often for multiple purposes. Tealeaf also supports proprietary file formats. The table below identifies the most common file formats in the Tealeaf system.

Note: The Search links below execute a search of Online Help for the file extension.

Tealeaf Proprietary Formats

A list of Tealeaf Proprietary Formats

Table 42. Tealeaf Proprietary Formats

Extension	Description and Usage	Search Link
.tla	Tealeaf Session Archive. Generated by Archive Session Agent or RTV.	https://community.tealeaf.com/dosearchsite.action?searchQuery.spaceKey=&searchQuery.queryString=.tla
.tlc	Tealeaf Session (Canister)	https://community.tealeaf.com/dosearchsite.action?searchQuery.spaceKey=&searchQuery.queryString=.tlc
.tle	Tealeaf event definitions. See "TEM Import-Export Tabs" in the *IBM Tealeaf Event Manager Manual*.	https://community.tealeaf.com/dosearchsite.action?searchQuery.spaceKey=&searchQuery.queryString=.tle
.tli	Tealeaf Image Archive	https://community.tealeaf.com/dosearchsite.action?searchQuery.spaceKey=&searchQuery.queryString=.tli
.tls	Tealeaf Session from RTV	https://community.tealeaf.com/dosearchsite.action?searchQuery.spaceKey=&searchQuery.queryString=.tls
.tlt	Tealeaf Report Template • Release 7.1 or later	https://community.tealeaf.com/dosearchsite.action?searchQuery.spaceKey=&searchQuery.queryString=.tlt
.tlx	Tealeaf Session List	https://community.tealeaf.com/dosearchsite.action?searchQuery.spaceKey=&searchQuery.queryString=.tlx
.tms	Tealeaf Management System Table File • Release 7.0 or later	https://community.tealeaf.com/dosearchsite.action?searchQuery.spaceKey=&searchQuery.queryString=.tms

© Copyright IBM Corp. 1999, 2014

Table 42. Tealeaf Proprietary Formats (continued)

Extension	Description and Usage	Search Link
.tmsv	Tealeaf Management System Index File • Release 7.0 or later	https://community.tealeaf.com/dosearchsite.action?searchQuery.spaceKey=&searchQuery.queryString=.tmsv
.tmsx	Tealeaf Management System Variable-Length Field File • Release 7.0 or later	https://community.tealeaf.com/dosearchsite.action?searchQuery.spaceKey=&searchQuery.queryString=.tmsx
.tsp	Tealeaf Spool File	https://community.tealeaf.com/dosearchsite.action?searchQuery.spaceKey=&searchQuery.queryString=.tsp

Other Common Formats in Use

Table 43. Other Common Formats in Use

Extension	Description and Usage
.cfg	Windows configuration file. Used by multiple components.
.csv	Comma-Separated Values file. Used for import/export.
.htm	HTML file. Alerts and reports can be exported in HTML format.
.ini	Windows Initialization file. Used by command-line tools.
.pdf	Adobe Portable Document Format file. Used for emailed reports and IBM Tealeaf cxVerify.
.xls	Excel file. Used for import/export.
.xml	eXtensible Markup Language file. Various uses.

Tealeaf Encryption Standards

Data security is a critical component of the Tealeaf solution. Tealeaf components employ the following encryption standards to safeguard your customers' interactions with your web application.

Key Characteristics

- Except as noted below, Tealeaf components use symmetrical encryption. The same key is used for encryption and decryption operations.
- Custom object identifiers (OIDs) are not used.
- Subject Alternate Names (SANs) are not used.
- SSL library is used to manage all encryption.

Encryption by Component

Component
Encryption

IBM Tealeaf CX Passive Capture Application Server
Web Server Keys are converted into encrypted 160-bit 3DES Tealeaf ptl formatted keys. Keys cannot be copied or used by another PCA server.

Data Encryption
For PCA servers pipeline privacy component or the Privacy pipeline agent, any part of the captured data can be encrypted to use either the RC2 or 3DES algorithm.

Server communication
By default, the supplied Tealeaf self-signed certificate/key is 2048-bit. Communication between the PCA server and the IBM Tealeaf CX Server can be encrypted with this certificate/key.

Note: SSL traffic between the PCA server and the IBM Tealeaf CX Server utilizes asymmetrical encryption.

Session Data
The Processing Server (Canister) stores all session data to disk in LSSN files encrypted using 3DES.

Note: Tealeaf maintains the 3DES keys, which are not configurable through Tealeaf products.

Portal and Report Server
Portal user passwords are encrypted using 192-bit 3DES.

Reference

- For more information on data privacy, see "Managing Data Privacy in Tealeaf CX" in the *IBM Tealeaf CX Installation Manual*.
- For more information on security, see “Implementing Tealeaf in a Secure Environment” on page 356.
- For more information on Portal authentication, see “Authentication” on page 329.
- For more information on SSL used in the IBM Tealeaf CX Passive Capture Application, see "SSL Keys" in the *IBM Tealeaf Passive Capture Application Manual*.

Chapter 25. IBM Tealeaf documentation and help

IBM Tealeaf provides documentation and help for users, developers, and administrators.

Viewing product documentation

All IBM Tealeaf product documentation is available at the following website:

https://tealeaf.support.ibmcloud.com/

Use the information in the following table to view the product documentation for IBM Tealeaf:

Table 44. Getting help

To view...	**Do this...**
Product documentation	On the IBM Tealeaf portal, go to **?** > **Product Documentation**.
Help for a page on the IBM Tealeaf Portal	On the IBM Tealeaf portal, go to **?** > **Help for This Page**.

Available documents for IBM Tealeaf products

Use the following table to view a list of available documents for all IBM Tealeaf products:

Table 45. Available documentation for IBM Tealeaf products

IBM Tealeaf products	**Available documents**
IBM Tealeaf CX	• *IBM Tealeaf Customer Experience Overview Guide* • *IBM Tealeaf CX Client Framework Data Integration Guide* • *IBM Tealeaf CX Configuration Manual* • *IBM Tealeaf CX Cookie Injector Manual* • *IBM Tealeaf CX Databases Guide* • *IBM Tealeaf CX Event Manager Manual* • *IBM Tealeaf CX Glossary* • *IBM Tealeaf CX Installation Manual* • *IBM Tealeaf CX PCA Manual* • *IBM Tealeaf CX PCA Release Notes*

© Copyright IBM Corp. 1999, 2014

Table 45. Available documentation for IBM Tealeaf products (continued)

IBM Tealeaf products	Available documents
IBM Tealeaf CX	• *IBM Tealeaf CX RealiTea Viewer Client Side Capture Manual* • *IBM Tealeaf CX RealiTea Viewer User Manual* • *IBM Tealeaf CX Release Notes* • *IBM Tealeaf CX Release Upgrade Manual* • *IBM Tealeaf CX Support Troubleshooting FAQ* • *IBM Tealeaf CX Troubleshooting Guide* • *IBM Tealeaf CX UI Capture j2 Guide* • *IBM Tealeaf CX UI Capture j2 Release Notes*
IBM Tealeaf cxImpact	• *IBM Tealeaf cxImpact Administration Manual* • *IBM Tealeaf cxImpact User Manual* • *IBM Tealeaf cxImpact Reporting Guide*
IBM Tealeaf cxConnect	• *IBM Tealeaf cxConnect for Data Analysis Administration Manual* • *IBM Tealeaf cxConnect for Voice of Customer Administration Manual* • *IBM Tealeaf cxConnect for Web Analytics Administration Manual*
IBM Tealeaf cxOverstat	*IBM Tealeaf cxOverstat User Manual*
IBM Tealeaf cxReveal	• *IBM Tealeaf cxReveal Administration Manual* • *IBM Tealeaf cxReveal API Guide* • *IBM Tealeaf cxReveal User Manual*
IBM Tealeaf cxVerify	*IBM Tealeaf cxVerify Administration Manual*
IBM Tealeaf cxView	*IBM Tealeaf cxView User Manual*
IBM Tealeaf CX Mobile	• *IBM Tealeaf CX Mobile Android Logging Framework Guide* • *IBM Tealeaf Android Logging Framework Release Notes* • *IBM Tealeaf CX Mobile Administration Manual* • *IBM Tealeaf CX Mobile User Manual* • *IBM Tealeaf CX Mobile iOS Logging Framework Guide* • *IBM Tealeaf iOS Logging Framework Release Notes*

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

© Copyright IBM Corp. 1999, 2014

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Bay Area Lab
1001 E Hillsdale Boulevard
Foster City, California 94404
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample

programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml.

Privacy Policy Considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. A cookie is a piece of data that a web site can send to your browser, which may then be stored on your computer as a tag that identifies your computer. In many cases, no personal information is collected by these cookies. If a Software Offering you are using enables you to collect personal information through cookies and similar technologies, we inform you about the specifics below.

Depending upon the configurations deployed, this Software Offering may use session and persistent cookies that collect each user's user name, and other personal information for purposes of session management, enhanced user usability, or other usage tracking or functional purposes. These cookies can be disabled, but disabling them will also eliminate the functionality they enable.

Various jurisdictions regulate the collection of personal information through cookies and similar technologies. If the configurations deployed for this Software Offering provide you as customer the ability to collect personal information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for providing notice and consent where appropriate.

IBM requires that Clients (1) provide a clear and conspicuous link to Customer's website terms of use (e.g. privacy policy) which includes a link to IBM's and Client's data collection and use practices, (2) notify that cookies and clear gifs/web beacons are being placed on the visitor's computer by IBM on the Client's behalf along with an explanation of the purpose of such technology, and (3) to the extent required by law, obtain consent from website visitors prior to the placement of cookies and clear gifs/web beacons placed by Client or IBM on Client's behalf on website visitor's devices

For more information about the use of various technologies, including cookies, for these purposes, See IBM's Online Privacy Statement at: http://www.ibm.com/privacy/details/us/en section entitled "Cookies, Web Beacons and Other Technologies."

Printed in USA